Are Feature Hierarchies Autosegmental Hierarchies?*

Colin Phillips, Massachusetts Institute of Technology

1. Introduction: On Division of Labour

Phonological representations are blamed for too much. They are taken to task for missing generalizations which can be captured elsewhere. In this paper I review some recent criticisms of the view that feature hierarchies are hierarchies of autosegmental tiers, especially the arguments of Hayes (1990) and Halle (1993), and try to show that although the considerations they raise are important ones, they do not force us to give up the standard view of feature structure. The criticisms can be more easily accommodated by making changes somewhere other than the theory of phonological representations.

Phonological theories capture generalizations about phonological processes in two ways. First, in the theory of representations; second, in the theory of transformations or rules — the theory of how different stages in a derivation take place. Generalizations about linguistic sound patterns can also be captured by phonetic theories of articulation and perception. Clearly, no subdomain has priority for accounting for any set of facts. I explore the trade-off between these different compoments of grammar in this paper.

1.1 Autosegmental Hierarchies

A great deal of work in phonology has attempted formalize the intuition that some groupings of features are natural classes and other groupings are not. Only the natural classes may participate in phonological processes. The idea that certain sets of features behave as natural classes is an old one (Trubetzkoy 1939), but closer examination of the idea was made vastly more feasible by the definition of the problem as one of finding a universal feature tree or feature hierarchy (Clements 1985; Mascaró 1983; Mohanan 1983; Sagey 1986). The assumption that (i) there are restrictions on the possible sets of features which may participate in phonological rules, and (ii) that these groupings are best expressed in terms of a universal hierarchical organization of features, is common to all approaches that I consider here.

* For advice, suggestions, criticisms and encouragement in the writing of this paper I am greatly indebted to Morris Halle. I would also like thank Andrew Carnie, Heidi Harley, Michael Kenstowicz, Shawn Kern, Jay Keyser, Donca Steriade, Hubert Truckenbrodt and Andrea Zukowski for useful discussion of this material. All errors are entirely my own.

MIT Working Papers in Linguistics 21, 173-226
Papers on Phonology and Morphology. © *1994. Colin Phillips*

1. *Hierarchical Organization*
   a. Only certain sets of features may participate in phonological processes.
   b. Hierarchical subgroupings of features define the possible classes of features.[1]
   c. The hierarchical organization of features is universally determined.

(2) illustrates the feature hierarchy which I assume in this paper[2]. However, my main aim here is not to justify the specifics of this hierarchy, although some details will turn out to be crucial in specific cases.

2.

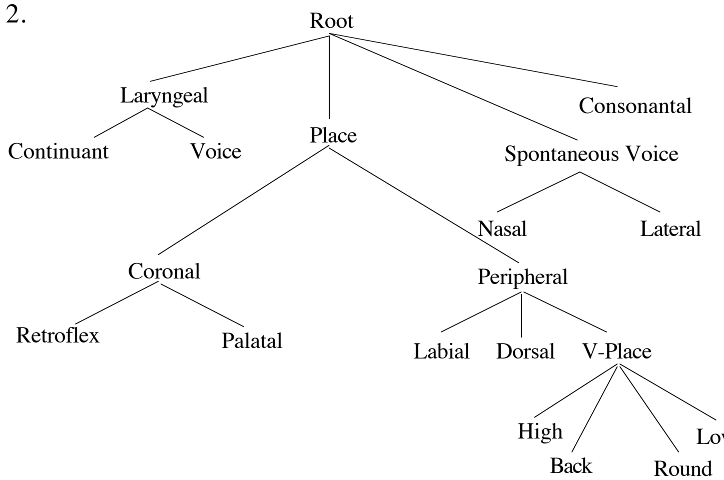

A more contentious issue is what hierarchies like (2) represent. The most common assumption is that each node in the hierarchy represents a unit on an autosegmental tier (3). Such autosegmental theories of feature structure generally share the following assumptions:

3. *Autosegmental Structure*
   All nodes in feature hierarchies are units on autosegmental tiers.

---

[1]*A priori*, this is by no means the only way of defining possible subgroupings — the alternatives are as varied as the possibilities for capturing regularities in natural language syntax. For example, the features may be assigned sequentially numbered indices, and only features with adjacent indices may participate together in phonological processes. Imagine the following hypothetical yet easily imaginable finding about place assimilation rules: rules exist which spread [labial] and [coronal] consonants, or [coronal] and [dorsal] consonants, or all three, but never just [labial] and [dorsal]. Such a generalization would not be expressible in terms of a hierarchical feature structure, although a restrictive framework like the one just suggested could express this.

[2]The hierarchy in (2) is based on Sagey (1986,1987) and incorporates various modifications proposed later. For motivation of [spontaneous voice] see Rice & Avery 1989; for [peripheral] see Rice & Avery 1991, Halle 1993; for V-Place see Sagey 1987, Clements 1991, Ní Chiosáin & Padgett 1993.

If we also assume that feature hierarchies are hierarchies of autosegmental tiers (4), then (5) follows as an automatic consequence.

4. *Autosegmental Hierarchies*
   Feature groupings are determined by the hierarchical associations between different autosegmental tiers.

5. *Node Sharing*
   Sets of features (i.e. subtrees) may be 'shared' by two segments by means of double linking.

Assumptions (3-5) are the core *representational* assumptions of autosegmental theories of feature structure, and (6) is the basic *transformational* assumption of these approaches. Section 2 discusses a problem for this set of assumptions, which can be resolved either by modifying (3-5) or by modifying (6).

6. *Constituent Spreading*
   a. Phonological rules operate on single nodes only.
   b. Phonological rules perform single operations only. (Clements & Hume, to appear, 7)

1.2 'Bottlebrush' Theories

While the assumptions in (3) to (6) have been the dominant ones, a growing number of phonologists have reached the conclusion that feature hierarchies should not be regarded not as autosegmental hierarchies, but simply as representations of category membership, which could equally well be expressed as labelled bracketings. In these approaches, only the terminal features of (2) retain autosegmental properties, so that what autosegmental structure is left has the appearance of what Hayes (1990) endearingly refers to as a 'bottlebrush', in which all terminal features are *directly* associated to the prosodic tier (7).

7.

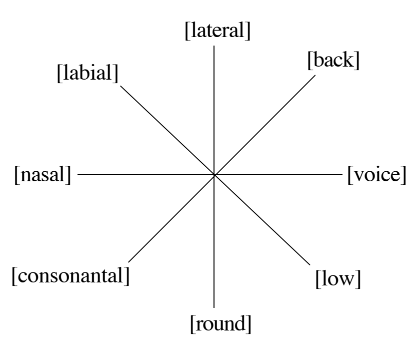

Hierarchical groupings of features can be superimposed on bottlebrush structures as in (8).

8.

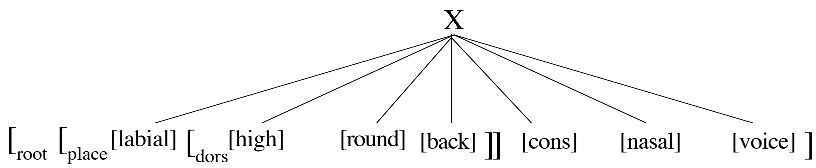

We should ask, what does the autosegmental view of feature structure commit us to that the bottlebrush view, augmented with feature groupings as in (8), does not? The primary difference is given as assumption (5), that class nodes can be shared by two or more segments. Class nodes cannot be shared in bottlebrush representations. Section 3 discusses evidence for class node sharing.

Second, class nodes do not encode timing information in an augmented Bottlebrush structure like (8), whereas they do under the autosegmental view. Association lines between autosegmental tiers have been taken to encode temporal overlap or simultaneity of the associated nodes (cf. Sagey 1988, Hammond 1988 for discussion). In autosegmental representations like (2), the temporal relationship between the skeletal tier and terminal features is mediated by temporal relations between class nodes. In representations like (8), on the other hand, temporal coordination of terminal features with the skeleton is unmediated. A consequence of (8) is that subgroups of features are not expected to be temporally coordinated: for example, if different subgoups of the features of a segment are realized at different times, the phonological representations do not predict that these subgroupings will correspond to class nodes in the feature hierarchy[3].

1.3 An Asymmetry in Tier Content

In presenting the received autosegmental view of feature hierarchies, so far I have been highlighting the parallels which this approach claims to hold between feature structure and the classic autosegmental domain — tone. I have concentrated on the interpretation of the association lines between autosegmental tiers. When we look at standard assumptions about the *content* of these autosegmental tiers, however, we find striking differences between what is

[3]I do not claim that it is impossible to build such predictions into an augmented Bottlebrush theory - Hayes (1990) does just that. What is important is that such predictions are automatic given autosegmental representations, whereas theories which assume representations like (8) need to add *extra* representational machinery to model group timing effects.

Of course, showing that group timing effects exist does not automatically favour the autosegmental approach: it needs to be shown that these effects must be encoded in the phonology, either because phonological processes refer to the temporal relations, or because they cannot be predicted on phonetic/articulatory grounds.

assumed about tonal tiers and what is assumed about the tiers of feature structure.

A typical assumption about the tonal tier is that it consists of sequences of units representing high and low tones. These may be associated to one or many skeletal units, or to none at all (9).

9.

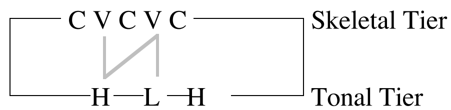

The Obligatory Contour Principle (10) imposes restrictions on autosegmental tiers containing nodes of more than one type. Its requirement that adjacent nodes be different amounts to the prediction that if, say, two adjacent syllables bear high tones, they are in fact realizing the same doubly linked unit on the tonal tier.

10. *Obligatory Contour Principle* (Goldsmith 1976)
Adjacent identical elements are prohibited.

In the formulation given in (10), the OCP is irrelevant to most of the tiers in an autosegmental feature hierarchy, because they only contain one type of unit. Only the terminal tiers representing '+' and '-' values of binary features have the content necessary for a condition like the OCP to apply. On the tiers containing just one node type, like [place], or [coronal], the OCP is either uniformly violated or inapplicable.

Nevertheless, phonologists have found evidence for OCP-like effects in feature structure at terminal and non-terminal levels. McCarthy (1988) cites restrictions on Ponapean consonant clusters (Itô 1986) as an OCP effect on [place], and suggests that co-occurrence restrictions on consonants in root morphemes of Proto Indo-European be analyzed as an OCP effect on [laryngeal].

By attributing these phenomena to the effects of the OCP on the [place] tier, it is implicitly assumed that the content of the [place] tier is more like the multi-valued content of the tonal tier than a series of identical units. In other words, it is assumed that it is possible to refer to variables picking out *values* of place, i.e. [αPlace] just as it is possible to refer to variables picking out values on terminal tiers, such as [αback].

By allowing well-formedness conditions to refer to values of both terminal and non-terminal nodes, all tiers are taken to have non-uniform content, like the tonal tier. Meanwhile, theories of phonological transformations have typically assumed a distinction between the tonal tier and terminal feature tiers on the one hand, where rules may refer to values, and class nodes on the other hand, where rules may not refer to values. If we eliminate this asymmetry, and allow rules to refer to variables at any level in feature hierarchies, we can account for some harmony processes which are problematic under standard assumptions about possible harmony rules, as I show in section 2. This modification also allows us to handle in an autosegmental framework phenomena which Halle argues require a 'Bottlebrush'-style analysis.

A harmony rule referring to values of a non-terminal tier can be expressed as in (11).

11.

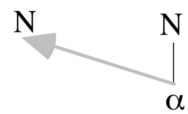

Spread αN leftwards

By writing rules referring to value variables like [αN] in (11) we allow two kinds of harmony processes which are unexpected under theories which do not allow reference to values at all levels.

Marked Values

First, we provide a new way of referring to the structurally more complex members of a natural class of features. For example, some coronal consonants are assumed to have extra structure below the coronal tier, whereas plain coronals are often assumed to have a [coronal] node with no dependent structure. Rules which pick out only those coronals with structure below the coronal tier may be construed as rules referring to [αCoronal].

In theories which do not allow reference to values at all levels, though, the only way to pick out the structurally more complex members of a natural class of segments is either (i) assume underspecification of the less complex segments: for the case of the coronals, this would mean that plain coronals would lack a [coronal] node; or (ii) allow rules to simply ignore segments which meet the structural description of the rule but are unmarked/structurally simple (Calabrese 1993; Carnie, this volume). In section 2.1 I present an example from Sanskrit which I claim requires the variable spreading account.

Multiple Values

In cases where a tier has more than one dependent in the same segment, the variable in a rule like (11) can range over more than one value, with the effect that a harmony rule will be implemented as more than one spreading operation, as illustrated in (12).

12.

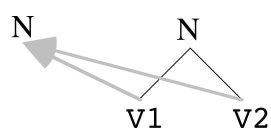

This implementation of single harmony rules as multiple operations entails weakening the Constituent Spreading hypothesis in (6) above, so that (6a) still holds, and rules still refer to single nodes, but (6b) is dropped, allowing rules to be *implemented* as multiple operations[4]. When rules refer to value variables, we make different predictions about rule blocking from Constituent Spreading approaches. It predicts that spreading of each value must be blocked individually, and, therefore, that harmony rules may be 'partially' blocked. I present a case from Scots Gaelic which motivates this assumption in section 2.2.

The rest of the paper is organized as follows: section 2 presents evidence for harmony rules which refer to value variables, and argues that this approach can accommodate one class of Halle's objections to the autosegmental view of feature hierarchies. Section 3 addresses Halle's second objection to autosegmental feature hierarchies — the claim that they are unneccesary — and argues that a variety of assimilation and dissimilation processes motivate the class node sharing that autosegmental theories predict.

Section 4 turns to Hayes' arguments that autosegmental feature hierarchies are inadequate for accounting for diphthongization phenomena. I argue that the range of attested diphthongizations does not justify Hayes' conclusions, with the result that diphthongization does not constitute an argument against autosegmental feature structures. Section 4 includes discussions of Icelandic preaspiration and of how contour segments should be represented.

2. Spreading Values

2.1 Sanskrit N‡ati

Assimilation processes are typically assumed to propagate a node rightwards or leftwards until another node of the same type is encountered. Normally, the blocking node occupies the same autosegmental tier as the node which spreads, and so the classes of 'spreaders' and 'blockers' are identical. The well-known N‡ati rule of Sanskrit (Whitney 1889; Schein & Steriade 1986; Cho 1991; Halle 1993) departs from this pattern: the class of spreading segments is a subset of the class of blocking segments. This leads to difficulties which appear to have gone unnoticed in previous analyses of N‡ati. (13) shows the consonantal inventory of Sanskrit.

13.

		Labial	Dental	Retr	Palatal	Velar	Laryngeal
	Stops	p	t	t.	c	k	
		ph	th	t.h	ch	kh	
		b	d	d‡	j	g	
		bh	dh	d‡h	jh	gh	

[4]It should be stressed that since rules refer to value variables only as an option, constituent spreading of the traditional kind is still possible.

Nasals	m	n	n‡	**	
Glides	v	l	r	y	
Fricatives		s	s.	s′	h

The N‡ati rule spreads retroflexion from a continuant /s.,r/ onto all following coronal nasals (14a-d), unless any other coronal intervenes, in which case no assimilation takes place (14e-h).

14.	a.	is. - n‡a	*seek*
	b.	pr - n‡a	*fill*
	c.	vrk - n‡a	*cut up*
	d.	ks.obh - an‡a	*quake*
	e.	bhug - na	*bend*
	f.	mrd - na	*be gracious*
	g.	marj - ana	*wiping*
	h.	ks.ved - ana	*hum*

I assume that what distinguishes the 3 series of coronals in Sanskrit is differences in their sub-coronal structure. The different series are characterized by either monovalent dependents of the coronal node, such as [retroflex] or [palatal], or by differing values of binary features like [anterior] and [distributed]. The difficulty with N‡ati is this: the class of harmonizing segments are picked out by their specific sub-coronal structure — the N‡ati rule must refer to this in its structural description. However, what characterizes the class of potential blockers of harmony is the presence of a [coronal] node.

(15) shows the formulation of N‡ati proposed by Schein & Steriade (1986) and Cho (1991). Since N‡ati assimilates two dependents of [coronal], the [coronal] node itself is assumed to spread rightwards. Therefore, propagation is blocked by *any* segment with a coronal node, due to the Line Crossing Constraint. This formulation successfully captures the range of blocking segments, but the way in which the class of spreading segments is picked out in (15) is problematic. (15) correctly states that [coronal] spreads only in certain contexts, but the context that it specifies is an unusual one[5].

[5]Halle (1993) expresses the rule in essentially the same way, as a rule spreading [coronal] when it dominates [-anterior,-distributed]. The only difference is that in Halle's framework the rule is implemented as separate spreading of [-anterior] and [-distributed]. For Halle the fact that all coronals block N‡ati follows from the fact that *all* coronals bear some specification for [±anterior] and [±distributed].

15.

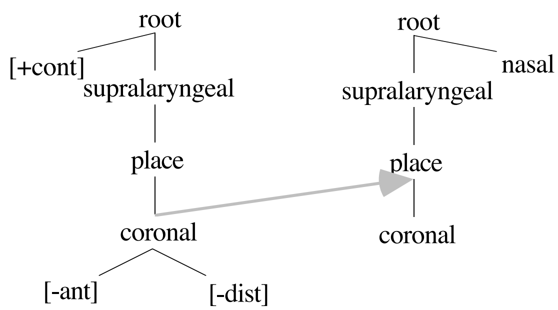

Typically, the node that harmonizes, and any extra context which must be satisfied for harmony to apply, are structurally disjoint, in the sense that neither node dominates the other. The [+continuant] and [+nasal] features specified as contexts in (15) are typical in this regard. But in addition to these, part of the context — [-anterior] and [-distributed] — is dominated by the spreading [coronal] node. If this kind of contextual specification is available, we run the risk of giving up one of the original motivations for postulating feature hierarchies, the prediction of what sets of features are natural classes. This is because by choosing an organizing node, eg, [root], and a subset of its dependents, it is possible to pick out any set of features as a 'natural class'[6].

If on the other hand we assume that monovalent features like [retroflex] and [palatal] distinguish the different series of coronals, it becomes straightforward to pick out the harmonizing segments /r/ and /s./. In this case, N‡ati spreads the [retroflex] node rightwards in continuant segments (16).

16.

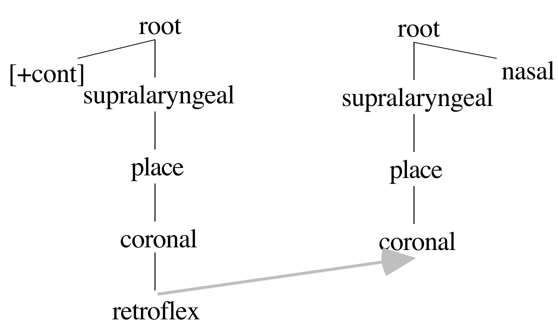

---

[6]If spreading rules may have 'dependent contexts', then *rules* can refer to arbitrary sets of features.

However, this does not automatically entail the more troubling prediction that arbitrary sets of features will be able to harmonize, if we assume that the set of features that spread are always exhaustively dominated by a single node, because there will presumably be some imaginable sets of features which are never exhaustively dominated by a single node in any segment, at any stage in a derivation. For this reason it is not clear to what extent the class of potentially harmonizing feature sets is enlarged by allowing dependent contexts.

Now, however, it becomes more difficult to capture the blocking of N‡ati. It's certainly not the case that all coronal segments are linked to the [retroflex] tier, so blocking cannot be a result of the Line Crossing Constraint. It's not clear that a locality condition like Archangeli & Pulleyblank's (1994) *Precedence Principle* helps either. This constraint rules out instances of gapped configurations like (17) in general.

17.

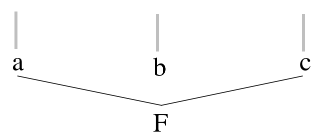

While this successfully prohibits assimilation in forms like (14e-h), this would also appear to prohibit well-formed instances of N‡ati like (14a-d), because consonant spreading across a vowel violates the *Precedence Principle*[7].

---

[7]Archangeli & Pulleyblank's *Precedence Principle* is intended as a constraint on association lines which both encompasses and is more restrictive than the Line Crossing Constraint. Its effects can be paraphrased as follows: branching nodes (like [F] in (17)) must have adjacent *anchors*, where anchors are moras for vowels and root nodes for consonants. In other words, there are strict locality requirements on multiply linked nodes, and adjacency relations are calculated at either the root or moraic level. See Archangeli & Pulleyblank (1994, ch. 1 & 4) for a more detailed exposition.

Vowels may spread across consonants in VCV sequences, since consonants are not linked to the moraic tier, which is the anchoring tier for vowels. However, a feature cannot spread from consonant to consonant in a CVC sequence, without also associating to the vowel. This is because vowels *are* represented on the anchoring tier for consonants. Although it is true that vowel harmony across consonants is vastly more common than consonant harmony across vowels, the effects of the Precedence Principle are too strong. Sanskrit N‡ati is a clear case of long distance consonant harmony. Shaw (1991) lists a number of other cases.

In addition, it is not clear that the effects of the Line Crossing Constraint follow from Archangeli & Pulleyblank's alternative. Since precedence relations for vowels are defined on a separate tier from consonants, *regressive* vowel harmony should be able to violate the LCC yet satisfy the Precedence Principle, as in the representation below.

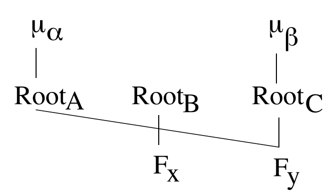

As far as I can tell, this representation is consistent with A&P's well-formedness conditions, given the stress on *precedence* rather than mere ordering: this requirement appears to have the effect that CV sequences but not VC sequences are subject to precedence conditions defined on the root tier. Nodes dependent on $Root_B$ (consonant) and $Root_C$ (vowel) preserve the linear ordering of $Root_B$ and $Root_C$. The

I suggest that these problems with the formulation of N‡ati can be avoided by using the variable spreading device introduced above. The N‡ati rule in (18) states that any *values* of continuant coronal segments are spread onto a following coronal nasal.

18.

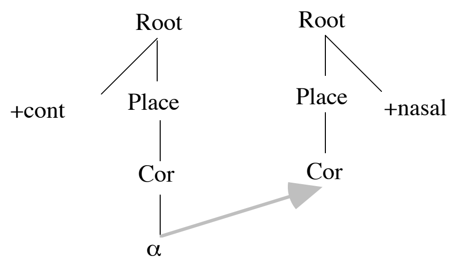

Let us see how this formulation manages to pick out the appropriate classes of harmonizing and blocking nodes. First, I assume that the 3 series of coronals have the representations in (19).

19.

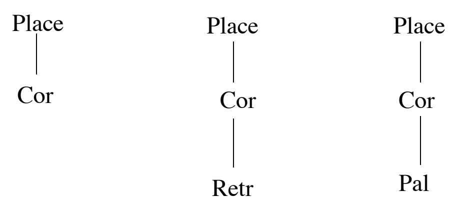

Plain coronals do not spread, since they lack a value of [coronal]. To see why the palatal coronals /y/ and /s'/ do not induce harmony, notice the gap marked by

---

line crossing involves nodes dependent on $Root_A$ and $Root_B$, which are not ordered by the *Precedence Principle*.

This makes the unexpected prediction that there is a linear asymmetry in vowel harmony processes. Crossed association lines matter for progressive harmony, so we expect to find blocking by intervening consonants; crossed association lines do not matter for regressive harmony, so there should be no blocking by selective consonants.

Surprisingly, this prediction seems to be fairly accurate cross-linguistically. The examples in section 2.2 and 3 below of vowel harmony processes blocked by selective consonants are overwhelmingly cases of progressive assimilation. The majority of instances of vowel harmony applying only across laryngeal consonants in Steriade (1987c) are also instances of progressive harmony. The only counterexamples are the complicated translaryngeal harmony rules of Wichita, Arbore and Yapese, which all impose additional restrictions on the V-features which can regressively assimilate, such that it is not clear that the restriction to intervening laryngeals only is due to the Line Crossing Constraint (vowel harmony in Fula pronouns (53) may fall into this category too).

Different construal of A&P's *Adjacency* condition might rule out the structure above, so that all LCC violations are ruled out; however, in the light of the preceding discussion, it might be best to allow some line crossings in.

stars in the palatal column of the inventory in (13) — Sanskrit lacks a palatal nasal[8]. If a [palatal] feature were to spread onto a coronal nasal, it would be blocked by Inventory Preservation.

The fact that all coronals block N‡ati is not surprising, given that the [coronal] tier is the one specified in the rule. A locality condition scans the tier specified by the rule, preventing targeting of the nasal in sequences like (20), since it is not adjacent to the trigger on the specified [coronal] tier.

20.

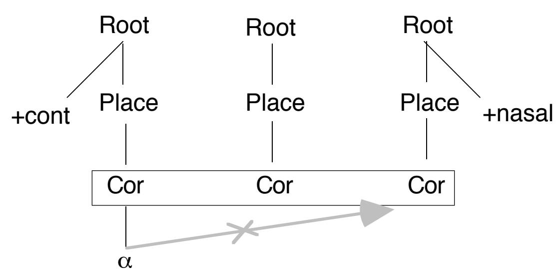

For Sanskrit, then, I exploit the variable spreading format to claim that *spreading* applies from the tier below [coronal] in N‡ati, whereas *blocking* applies on the [coronal] tier itself. The next section tries to show how the innovations introduced in this analysis of N‡ati helps to fend off some of Halle's arguments against the autosegmental view of feature hierarchies.

2.2 Barra Vowel Copy & Terminal Feature Spreading

By allowing variable spreading rules such as the one I have proposed for Sanskrit Nati, I am giving up the strong version of the Constituent Spreading hypothesis in (6a-b) above. Any constituent of a feature hierarchy may be spread by a rule as in Clements (1985), but I have dropped the assumption that rules may spread *only* single constituents. The weaker requirement that I impose is that when a set of nodes spread, they are exhaustively and immediately dominated by a single node in the feature hierarchy. In this section I compare this approach with an approach to the interpretation of feature structures recently proposed by Halle (1993) (see also Halle & Vaux 1994). Halle's alternative takes a very different view of what feature structures represent, but allows for very similar analyses of a number of phenomena.

---

[8]The assumption that Sanskrit lacks a palatal nasal departs from most analyses, which follow the inventories given by both Pan.ini and Whitney (1889). Emeneau (1946) points out that once foreign borrowings and inventions of the grammarians are excluded, the palatal nasal /ñ/ is found only before the palatal stops c, ch, and j, and after c and j. Emeneau concludes that the palatal nasal which appears in the *devanagari* orthography should be treated as a positional variant of the dental /n/. He suggests that palatal /ñ/ appeared in the grammars and orthographies due to the same desire for symmetry which led Sanskrit grammarians before Pan.ini to include a bogus long syllabic lateral /l⁄/ in the inventory.

Halle (1993):
Autosegmental hierarchies are
(i) unnecessary,
(ii) sometimes inadequate.

Halle (1993) assumes that organizing nodes in feature hierarchies do nothing more than represent groupings of features; they are not elements on autosegmental tiers. Individual terminal features, and these alone, exhibit autosegmental properties like multiple linking in harmony processes. Halle's view appears closer to Clements' proposal than it actually is, because Halle retains feature tree representations. However, he assumes that the trees only represent category membership information. Although not explicitly formulated as such, this view amounts to a version of the Bottlebrush theory (cf. (7) & (8) above), in which each feature is directly associated to the root[9]. Nevertheless, Halle assumes that harmony rules may be stated as we have become accustomed to, and may refer to organizing nodes, such as [place]. The only difference is in the *implementation* of rules: when a rule states that an organizing node must be spread, this is actually implemented as the spreading of all terminal nodes dominated by that organizing node, as in the complete vowel copy rule in (21).

21. Rule: Spread [place] rightwards

Implementation:

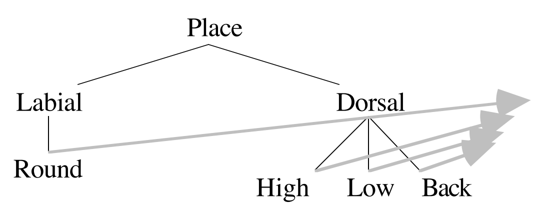

This is a stronger modification of the Constituent Spreading hypothesis than I have proposed. Whereas my variable spreading rules require that sets of spreading nodes be exhaustively and immediately dominated by an organizing node, Halle merely requires that sets of spreading nodes be exhaustively dominated by an organizing node — there is no immediate domination requirement. In addition, non-terminal constituents *cannot* spread for Halle[10].

One of the clearest predictions of Halle's approach is that, since harmony takes place on a feature-by-feature basis, the blocking of harmony

[9]One difference between Halle's view and the Bottlebrush view as presented by Hayes is that Halle assumes that feature tiers all associate to a Root tier rather than directly to the skeletal tier. This allows him to carry over into a Bottlebrush approach McCarthy's (1988) observation that the features [±sonorant] and [±consonantal] only ever assimilate when all features assimilate. For the purposes of this paper, however, Halle's view may be regarded as equivalent to the Bottlebrush theory in (7-8).
[10]The exceptions are the features [consonantal] and [sonorant], which Halle assumes to be the root node, following McCarthy (1988).

processes can only occur in a similar feature-by-feature fashion. So, for example, the vowel feature [+high] should never be blocked from spreading by a segment which lacks specification for [αhigh]. The fact that a pair of segments share an organizing node should never have a bearing on how they interact in harmony processes — only terminal features are relevant for blocking. On the other hand, the view I have proposed does allow segments to share organizing nodes. In the next few sections I present some cases which motivate the terminal feature spreading assumption and show that they are equally consistent with the variable spreading view I've suggested. I also present examples of assimilation and dissimilation which seem to require analyses involving organizing node sharing.

Halle's approach and my modified version of the autosegmental view both allow a single harmony process to be implemented as a number of distinct association operations, of which each may be blocked individually, with the result that 'partial harmony' effects are predicted to arise. The paradigm example of this is the vowel copy process found in the Barra dialect of Scots Gaelic (Borgstrøm 1937, 1940; Clements 1986; Sagey 1987; Halle 1993). The inventories in (22-23) show that among the features distinguished by segments in Barra, [±back] is distinctive for both vowels and consonants. This leads to an interesting interaction when it comes to determining the quality of an epenthetic vowel. Epenthetic vowels are identical in quality to the preceding vowel (24), unless the intervening consonant conflicts in backness, in which case, the vowel has the backness value of the intervening consonant (25).

22.

	i	„	u	e	ˆ	o	æ	a	¿	
	+	+	+	-	-	-	-	-	-	High
				-	-	-	+	+	+	Low
	-	+	+	-	+	+	-	+	+	Back
		-	+		-	+		-	+	Round

23.

	Plain			Palatal ([-back])		
p	t	k	-	c&	k'
b	f	g	-	j&	g'
f	s	¯	-	s&	ç
v		'	-		j
m	n, N		-	-, N'	
	r, R			r', -	
	(l), L			l', L'	

24. a. /u/ duN**u**x\ ' *Duncan*
    b. /„/ „n**„**\s *Angus*
    c. /æ/ m**æ**s&ir' *time*

25. a. u → i bul'ik' *bellows (gen. sg.)*
    b. „ → i d„r'i *fishing line (gen. sg.)*

c. æ → a s&ærav *bitter*

As Sagey (1987) points out, the fact that all vowel features are copied to the epenthetic vowel in (24) would normally lead us to assume a rule spreading a node which dominates all vowel features, eg. [place] or [V-Place]. But when a conflicting value of [back] intervenes between the harmonizing vowels, as in (25) a Constituent Spreading analysis does not predict *any* vowel features to harmonize. The fact that only [back] is blocked is expected, though, if vowel copy is implemented as multiple operations, due either to a variable spreading rule (26a) or due to a terminal spreading interpretation of feature hierarchies, as in (26b).

26a.

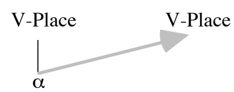

b.

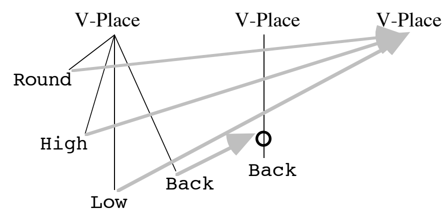

We should counter a couple of alternative analyses, which could claim that the Barra facts are consistent with maintaining a strict version of the constituent spreading hypothesis.

The first possible objection is that feature-by-feature blocking of harmony is an illusion: the one feature which is claimed to be ‘blocked’ is [back], which is in fact *always* supplied by the intervening consonant. We need to find examples where the intervening consonant cannot supply a value for [back]. (24c) shows that backness values can in principle spread from the preceding vowel: /m/ is non-contrastive for backness, so the [-back] feature on the epenthetic vowel in (24c) must be due to the preceding vowel. This supports the claim that the feature is actually blocked from spreading when a non-labial consonant intervenes.

Another possible reanalysis of the Barra facts consistent with Constituent Spreading would be to claim that when [back] is blocked from spreading, the remaining features — [round], [low], and [high] — *do* still spread as a constituent. This requires a revision of the feature hierarchy in (2) such that [round] and the height features form a constituent to the exclusion of [back], as in (27).

27.

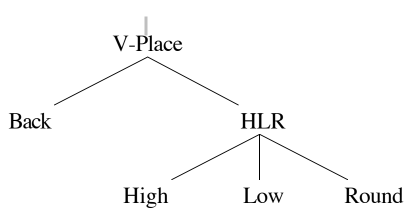

In this case Barra vowel copy might be explained as spreading a node [V-Place] which dominates all 4 features when no Line Crossing Constraint violation results, and spreading the lower node [HLR] dominating [high], [low], and [round] otherwise. This approach makes very strong predictions about the range of partial vowel harmony processes to be found across languages: the structure in (27) predicts that whenever [back] and at least one other vowel feature harmonizes, *all* other vowel features must harmonize. On the other hand, the analysis of Barra which claims that each feature spreads individually predicts that other subsets of vowel features may be assimilated, provided that appropriate blockers are present.

Evidence to distinguish these two positions is readily available. Odden (1991) uses data from Wikchamni (Gamble 1978; Archangeli 1985) as evidence for a node [Back-Round] in the feature hierarchy which exhaustively dominates [back] and [round]. Wikchamni has the 5 vowel system in (28). This language harmonizes [back] and [round] from left to right between vowels in consecutive syllables, provided that they have identical values for the feature [high].

28.

	i		u	o	a
high	+	+	+	-	-
back	-	+	+	+	+
round	-	-	+	+	-

29.

a.	pin’s&i	_	pin’s&i	*stung*
b.	huts&i	_	huts&u	*knew*
c.	tï iss&i	_	tï ïss&ï	*made*
d.	t^{h}ans&i	_	t^{h}ans&i	*went*
e.	tawt^{h}at	_	tawt^{h}at	*might run*
f.	t’oyxat	_	t’oyxot	*might doctor*
g.	hukyat	_	hukyat	*might mix*

Halle (1993) shows that the spreading of [back] and [round] alone can be readily analyzed in his terminal-feature spreading framework, thereby maintaining a feature hierarchy in which vowel features are dominated by the articulator nodes [labial] and [dorsal][11]. He assumes that identical height features fuse if they are identical (following Cole & Trigo 1988), so that the height harmony required for back-round harmony amounts to a requirement that the harmonizing vowels must share a [high] node. Then a rule harmonizing [place] (i.e. full vowel harmony) has the effect of spreading [back] and [round], since values of [high] are already identical.

[11]Retaining an articulatorily motivated feature hierarchy is one of Halle’s main goals in his paper, and it drives some of his conclusions about the nature of feature structure and possible rules. I have nothing to offer here on the question of whether phonological representations should match articulatory plans or the structure of the vocal apparatus.

More important than the specifics of the analysis, however, is the fact that Wikchamni shows precisely what a Constituent Spreading analysis of Barra predicts to be impossible — it spreads [back] together with another vowel feature, but not all other vowel features. The fact that it is impossible to analyze vowel harmony in both Barra and Wikchamni as single processes in the Constituent Spreading framework constitutes a strong argument in favour of the modification of assumptions (6a-b) above, as Halle and I propose[12].

Halle's conclusion from Barra vowel copy is that the constituent spreading hypothesis must be abandoned in favour of a Bottlebrush approach in which only terminal features spread. I have tried to show that Barra vowel copy and Sanskrit N‡ati can be handled by a theory which drops the strict constituent spreading hypothesis, but which maintains an autosegmental view of feature hierarchies.

Halle (1993) also presents terminal feature spreading analyses of a wide variety of examples used by Odden (1991) to argue for perceptually motivated subgroupings of vowel features. Halle claims that Odden's arguments for the partial feature hierarchy in (30) do not go through once a terminal feature spreading approach is adopted.

30.

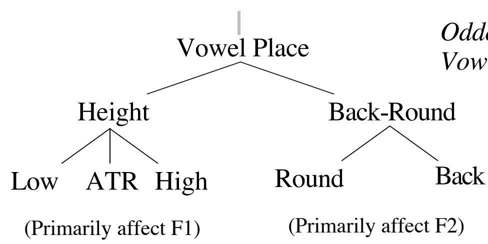

Odden's perceptually motivated Vowel Geometry

Although I cannot give a full treatment of all of Odden's examples here, it should be pointed out that if Odden's vowel geometry in (30) is correct, my variable-spreading analysis of Barra vowel copy is impossible, since my account relies on the assumption that there is a node which *immediately* dominates all terminal vowel features. My analysis therefore becomes untenable if there are subgroupings of vowel features in the feature hierarchy. Fortunately, Halle's reanalyses of Odden's examples using terminal feature spreading can be readily

[12]Ní Chiosáin (1994) proposes to reconcile the Barra epenthesis facts with the constituent spreading hypothesis by by attributing the quality of the epenthetic vowel to independent processes. She suggests that (i) only height features ever spread from the preceding vowel; (ii) backness is always determined by the consonants separated by epenthesis; (iii) roundness is always predictable. Assumption (ii) is consistent with all of the forms here and in Ní Chiosáin (1994), but the vowels separated by epenthesis are not always contrastive for [±back]. In examples in which *both* consonants are non-contrastive for [±back], the vowel copy analysis predicts that backness will be determined by the preceding vowel; Ní Chiosáin's account does not predict this. Unfortunately, I do not know of relevant examples. Assumption (iii) is only partially justified in Ní Chiosáin (1994).

reinterpreted as variable spreading analyses, so it does not seem necessary to introduce sub-groupings of vowel features.

This section has focussed on cases which Halle (1993) claims can be analyzed in his version of the Bottlebrush theory, but which present difficulties for autosegmental approaches. I've tried to show how the autosegmental approach can be modified to handle these problems. In the next section I turn to Halle's second, implicit, argument for the Bottlebrush theory: there are no phenomena which demand that class nodes in feature hierarchies be treated as autosegmental nodes.

3. Class Node Sharing

3.1 On Some Apparently Disjunctive Rules

Sagey (1987) recognized the difficulty of the Barra vowel copy facts for theories which assume Constituent Spreading. Sagey's analysis is essentially the same as Halle's, but she stops short of concluding that terminal feature spreading is how assimilation is implemented in all cases, due to other examples which seem to demand an analysis in terms of Constituent Spreading.

Ainu

Sagey presents vowel harmony in Ainu (Itô 1984) as a minimal contrast with Barra vowel copy. In this language the transitivizing verbal suffix is usually a copy of the root-final vowel (32). Vowel copy crosses any consonant unimpeded, with the exception of the glides /w,y/, in which case the suffix vowel is always /e/ (33). (31) shows the phonemic inventory of Ainu, taken from Maddieson (1984).

31.

p	t	c*	k	h
	s			
m	n			
	rr			

	i	u	e	o	a
high	+	+	-	-	-
low	-	-	-	-	+
back	-	+	-	+	

32.			33.		
a.	mak-a	*to open*	a.	ray-e	*to kill*
b.	ker-e	*to touch*	b.	chaw-e	*to solve*
c.	pis-i	*to ask*	c.	hew-e	*to slant*
d.	pop-o	*to boil*	d.	piw-e	*to cause to run*
e.	tus-u	*to shake*	e.	poy-e	*to mix*
			f.	huy-e	*to observe*

Sagey's analysis of the blocking effect of glides runs as follows. Vowel copy must spread the features [high], [low], and [back]; all of these features are blocked from spreading by the presence of intervening /w, y/; but since the glides do not have *all* features of vowels, a terminal feature spreading analysis like the one given to Barra vowel copy incorrectly predicts that glides only partially block vowel copy in Ainu. Sagey therefore concludes that Ainu vowel copy must be analyzed as spreading of the constituent dominating all vowel features ([2-Place]), in order to exclude the possibility of partial blocking.

Recall that Halle's approach only predicts full blocking if there is an intervening segment which matches the spread nodes feature-by-feature. Thus, Halle must assume that Ainu glides are marked for [high], [low] and [back]. This does not follow from any version of the distinctive specification approach to feature marking which Halle assumes, since the two glides contrast only in backness[13]. Therefore he necessarily assumes that the glides are positional variants of vowels, which ensures that they block spreading of all vowel features.

Even if Halle's conjecture about the underlying forms of glides in Ainu is correct, there is still reason to prefer the constituent spreading analysis over one in which features are spread individually. The argument is based on the analysis of the additional forms in (34). All of the suffixes in (34) are [+high], and bear the opposite value of [back] to the root vowel. Itô argues that whereas the suffixes in (32-33) are underlying floating V's with no features, the suffixes in (34) are underlyingly specified [+high]. In these cases the rule in (36) determines the backness value of the suffix:

34. a. hum-i *to chop up*
b. pok-i *to lower*
c. pir-u *to wipe*

35. a. kar-i *to rotate*
b. ram-u *to think*

[13]The assumption that only contrastive features are specified is important for Halle's approach.

Since he assumes that features are directly associated to the root node, locality conditions on rule application can only refer to the root tier or to the tiers which terminal features occupy: therefore, spreading rules can be constrained by strict string adjacency or by the Line Crossing Constraint, but nothing else. In particular, it is impossible to rely on adjacency conditions on class nodes.

Given these limited resources for expressing locality conditions on rule application, it is difficult to capture the contrast between coronal harmony rules like Sanskrit N‡ati and Tahltan coronal harmony (Shaw 1991). Recall that N‡ati spreads a subset of coronals, but all coronals block this harmony. In Tahltan, a subset of coronals harmonize, and the harmony is not blocked by any other segments. A straightforward way of capturing the contrast between the two rules would be to assume that N‡ati must apply to segments which are adjacent on the [coronal] tier, whereas Tahltan coronal harmony is unbounded.

Halle's framework cannot make this distinction, since class nodes are not autosegmental nodes; instead he derives the contrast from the structure of the inventories of Sanskrit and Tahltan. In both languages coronal harmony spreads the features [anterior] and [distributed]; in Sanskrit these features are contrastive for all coronals, but in Tahltan these features are only contrastive for the coronal series which harmonize. As a result, what looked like contrasting locality conditions in the two languages follows from the Line Crossing Constraint.

Given the interesting results Halle derives from contrastive specification for Sanskrit and Tahltan, it would be unfortunate to give up contrastive specification in order to capture the blocking effect of Ainu glides.

d. ket-u *to rub*

36. *Melodic Dissimilation Rule (MDR)* (Itô 1984, 508)
[+high] → [-αback] / [αback] ____

At first glance, we would not want this rule to apply to all representations, as it would exclude all of the cases of full vowel harmony in (32). Itô assumes that the rule applies *unless* vowel identity results from "autosegmental spreading of a single melodic unit" (1984, 512). In terms of more recent approaches to segmental structure we can say that dissimilation occurs unless the two vowels *share* all of their features: in this case the environment of the rule is not met, since there is no sequence of two nodes linked to [αback] which could be subject to the MDR[14]. In terms of hierarchical representations of features, the MDR is an OCP constraint requiring contrasting values of [back] when adjacent nodes immediately dominate [back]. When 2 segments share a node on the tier immediately dominating [back] (eg. [V-Place] or [Dorsal]) the MDR does not apply. (37a) shows an environment in which MDR applies, (37b) shows an environment where MDR fails to be triggered[15].

37a.

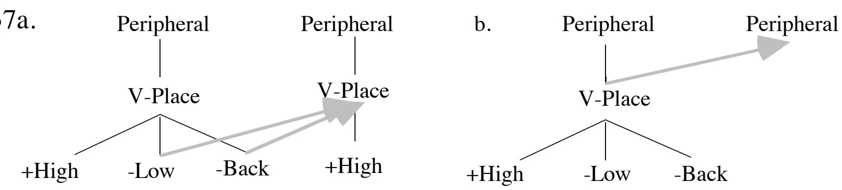

Notice that this analysis of where MDR does and does not apply crucially relies on the availability of class node sharing, which is unavailable in a Bottlebrush approach like Halle's.

Ngbaka

The scenario we have just observed in Ainu, in which adjacent vowels must bear contrasting values for a particular feature, unless they are identical in all features, is found in a number of other unrelated languages. In Ngbaka (Westcott 1965; Chomsky & Halle 1968) there is a constraint on the distribution of vowels in disyllabic words to the effect that the two vowels must either contrast in the value of [αhigh], or be identical in all features[16]. The vowel inventory of Ngbaka is shown in (38).

---

[14]Here it is important to distinguish *sharing* of features from *identity* of features, since backness dissimilation can apply to sequences of 2 vowels which have the same values of [high] and [low] (eg. 32c).

[15]Notice that the Ainu facts present an argument for a feature hierarchy in which vowel features are all dominated by a special class node like [V-Place], rather than by the major articulator nodes. If Ainu vowel copy were a rule spreading [dorsal], we would expect it to be blocked by velar consonants. (32a) shows that velars are transparent to vowel copy. [Note that this argument does not apply to theories in which terminal features spread, like Halle's.]

[16]/a/ is not subject to this restriction, so it may appear with any other vowel.

38. [+high] i u
[-high]$_{[-low]}$ e o
[-high]$_{[+low]}$ « a ¿

This means that, for example, /i/ cannot appear in a disyllabic word with the other [+high] /u/; /o/ cannot appear with [-high] /e, «, ¿/.

Tzeltal

Tzeltal (Slocum 1948; Kaufman 1971) has a number of VC suffixes in which V is either identical to the root vowel, or has the opposite value for [αback]. Tzeltal has the 5 vowel system in (39).

39.

[-back]		[+back]
i		u
e		o
	a	

40. -Vl - 'place where thing abounds'

a.	c&'umil	*squash-patch*
b.	pahc&'ul	*pineapple plants*
c.	c&enk'ul	*bean-patch*
d.	mayul	*tobacco plants*
e.	ic&il	*chile-patch*

41. -Vl - possessor (inanimate)

a.	swicul	*its hill*
b.	slewul	*its fat*
c.	spak'ul	*its cloth*
d.	slumil	*its ground*
e.	spos&il	*its medicine*
f.	sc&ayul	*its fish*
g.	si ul	*its firewood*
h.	ste el	*its wood*
i.	sti il	*its edge*

Guere

Finally, Guere (Paradis & Prunet 1989) has a restriction on the height of adjacent vowels, which is a slightly more complicated version of what we have already seen in Ngbaka. Ngbaka restricts the appearance of sequences of [αhigh][αhigh] vowels to instances of complete identity. In Guere, there are no sequences of [-high][-high] *unless* the vowels are (i) fully identical and (ii) either a coronal or nothing at all intervenes between them. (42) shows the vowel inventory of Guere.

42.

[+ATR]			[-ATR]	
i	u			‹
e	o	«		¿
			a	

(43-44) show that sequences of [αhigh][-αhigh] and [+high][+high] are permissible.

43. *High, nonhigh*

a. k ɪo *rat* b. n m« *bird*
c. k‹la *hand* d. ɪuo *round*

Nonhigh, high

e. ɪai *robe* f. gw« *burn*!
g. gbau *box* h. jreu *monkey*

44. *High, high*

a. nimi *animal* b. z g‹ *chameleon*
c. du u *chest* d. bui *ashes*

Nonhigh, nonhigh

e. —-

(45-46) show how [-high][-high] sequences are avoided: when a [+high] object pronoun suffix attaches to a stem ending in a [-high] vowel, the stem vowel remains constant (45). However, when a [-high] object-pronoun suffix is attached to a stem ending in [-high], the stem vowel changes from [-high] to [+high] (46).

45. a. n«« + ‹ → n«‹ *stick it!* (*n ‹)
b. gblee + ‹ → gble‹ *welcome it!* (*gbl ‹)
c. w¿(l)¿ + ‹ → w¿(l)¿*wash it!* (*w‹(l)‹)

46. a. n«« + « → n « *stick it!* (*n««)
b. gblee + « → gbl « *welcome it!* (*gble«)
c. w¿(l)¿ + ¿ → w‹(l)¿*wash it!* (*w¿(l)¿)

Paradis & Prunet assume that the alternations in (45-46) are due to an OCP constraint which rules out adjacent [-high] features. The exceptions to the height constraint are shown in (47-48): they occur when the 2 [-high] vowels are identical in all features, and either string-adjacent (47), or separated by a coronal (48).

47. a. baa *manioc* c. ¿¿ *week*
b. y«« *to dry*

48. a. w¿ ¿ *to wash* c. s¿ ¿ *to lose weight*
b. be e *to hang*

The apparently disjunctive conditions on vowel features in Ainu, Ngbaka, Tzeltal and Guere can all be straightforwardly expressed in a theory which assumes that class nodes occupy autosegmental tiers. In each language, there is a ban on adjacent identical values of a given feature, but the ban can be evaded if two vowels share a class node like [V-Place], because in this situation the two vowels only have one set of terminal vowel features. (50) illustrates this using the ban on [αhigh][αhigh] in Ngbaka. The other languages work in essentially the same way[17].

49.	Ainu	*[αback][αback] except under identity
	Tzeltal	*[αback][αback] except under identity
	Ngbaka	*[αhigh][αhigh] except under identity
	Guere	*[-high][-high] except under identity

50.

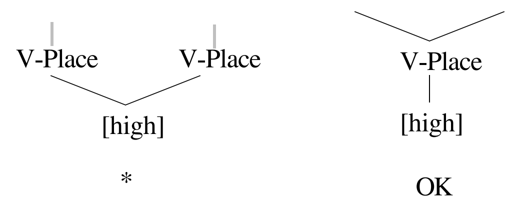

These apparently disjunctive conditions are not expected given just the formal properties of a bottlebrush view of feature structures, because only terminal features can be shared under this approach. This does not automatically favour the autosegmental view of feature hierarchies, though. We must consider the possibility of accounts of the disjunctive rules which do not rely on class-node sharing. It would be particularly interesting if the kinds of facts discussed in this section could be shown to follow from independent factors, so that they do not need to be captured by the properties of the formal representations chosen. Steriade (p.c.) suggests just such a possibility, whereby the facts of Ngbaka, Ainu, Guere and Tzeltal might follow from phonetic-perceptual factors.

Steriade suggests that the generalizations we have reviewed should be viewed rather differently. Let us take Ngbaka as an example: as I presented the facts, adjacent vowels *must contrast* in height, unless they can escape this requirement by being identical and linked at a non-terminal level. Steriade proposes to express the generalization in almost the opposite fashion: adjacent vowels *cannot contrast* in features other than height (e.g. [back]) unless they also contrast in height. Steriade suggests that this licensing of contrasts is plausibly grounded in perceptual facts: the presence of a height contrast between vowels, which is easily perceived, might improve the likelihood of perceiving other contrasts between the two vowels. (51a) is a schematic plot of how contrast in

[17]The Guere constraint requires a little more elaboration. First, the constraint only rules out adjacent minus values on the [high] tier. Paradis & Prunet (1989) suggest that this may be the result of underspecification of [+high] at the point in the derivation where the constraint applies. Second, sequences of identical [-high] vowels only escape the constraint if they are string adjacent or separated by a coronal. Paradis & Prunet claim that these are the only environments in which vowels can share a class node and thereby escape the height constraint. This could be either because coronals are underspecified, as Paradis & Prunet claim (cf. (55) below), or because the node which exhaustively dominates all vowels features in Guere is absent from even a fully specified coronal.

height (F1) between two vowels might aid the perception of contrast in backness (F2). (51b) illustrates how a contrast licensing rule could be expressed as a well-formedness condition on feature structures.

51a.

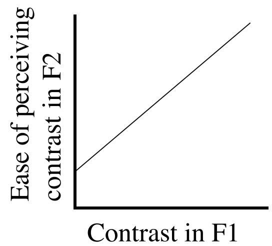

b.

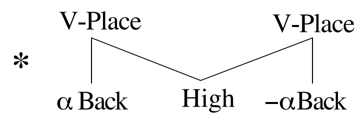

This amounts to an extension of the domain of what Steriade has called *Positional Neutralization* (Steriade 1993b,1994) from feature identification to feature discrimination. With respect to the identification of features, Steriade has argued that the correct identification of features by the listener is made easier in two ways. First, contrastive features may be neutralized in positions in which they are difficult to perceive; second, contrastive features which are hard to perceive may be extended in duration (assimilation). Extending the spirit of this proposal to discrimination of two segments, Steriade suggests that contrasts between two segments which are hard to perceive are only permitted in environments which aid their correct discrimination.

The two perspectives on these facts are quite different. The view I presented focusses on the prohibition of adjacent non-contrasting features, whereas Steriade's view emphasizes the role of some contrasts in making other contrasts more salient. The identical vowels which I claim 'escape' the OCP constraint are for Steriade precisely the vowels which are subject to Contrast Licensing constraints like (51b). Steriade's approach has the attraction of being both restrictive and phonetically grounded, but it faces a couple of difficulties.

First, Steriade's approach predicts that only contrasts license other contrasts. We do not expect that contrast in [back], [round], [ATR] should depend on sequences with *either* contrasting height *or* adjacent [+high][+high], which is what we find in Guere. Under the OCP view suggested here, which emphasizes the licensing of identity, the Guere situation, in which only adjacent [-high] segment prevent contrasts in other segments, is far less surprising: it is just a more specific variant of the constraint in (50), in which the OCP only applies to [-high].

Second, Steriade's view makes extremely strong predictions about the connection between phonological contrasts and the perception of contrasts. Classic findings about categorical perception show that listeners' ability to perceive contrasts between consonants closely matches the phonological significance of the contrast. i.e. 5ms differences in voicing onset time (VOT) between stops are very much more likely to be perceived when the 5ms spans a category boundary. Eimas (1963) found that vowel *identification* is very similar to consonant identification insofar as vowels with small acoustic differences are fairly easily mapped onto different categories. However, vowel *discrimination* is rather different from consonant discrimination: acoustic differences between vowels can be equally well discriminated whether they span a category boundary or not, whereas discrimination among consonants within a category is extremely difficult.

Recall that Steriade's hypothesis is that discrimination of a contrast along one dimension is aided by the additional presence of a category boundary along a second dimension. Since Ngbaka distinguishes three levels of vowel height, and only the difference between the highest vowels and all other heights is relevant to allowing contrast in [back] and [round], it would have to be category boundary between [±high] rather than direct acoustic differences which licenses other contrasts. However, if we assume based on Eimas' findings, that just as category boundaries do not enhance the perception of contrasts along their own acoustic dimension, category boundaries do not have any special status in facilitating the discrimination of contrasts on another dimension, then we expect the low-mid contrast (both [-high]) in Ngbaka to have the same facilitatory effect on [back] or [round] discrimination as the mid-high contrast (across boundary [±high]), since the acoustic contrast is equivalent. This is however not the case[18].

3.2 Consonantal Transparency Effects

Assimilation processes have guided our thinking about feature hierarchies in two different ways. The bulk of research has focussed on the problem of classifying the groups of features that are affected together in assimilation processes. These groups are used to define the organizing nodes in feature hierarchies. Any results of this line of inquiry are equally applicable to an autosegmental theory of feature hierarchies or a bottlebrush theory incorporating an independent notion of hierarchical grouping of features, such as Halle's or Hayes'.

Another kind of work on feature hierarchies focusses on what kinds of long-distance assimilation rules exist, and what kinds of elements they are blocked by. Examining structural similarities between spreading and blocking segments provides clues to what nodes are being assimilated. This kind of approach is more contingent on the specific assumptions of autosegmental feature hierarchy theories. These approaches assume that since autosegmental spreading can take place on non-terminal planes, it is sufficient for a pair of segments to have a class node in common for one of them to block spreading

[18]This argument could turn out to be invalid given more careful examination of the phonetic facts: the contrast in F1 between low and mid vowels could turn out to be very small compared to the contrast between mid and high vowels. Based on the English vowel formant data in Ladefoged (1975/1993) the mid-low contrast in F1 is slightly smaller than the mid-high contrast on an absolute kHz scale. This discrepancy is likely to be even greater in perceptual terms, since lower vowels have higher F1 and pitch discrimination is known to use a logarithmic scale.

from the other. Only where blocking can be reduced to blocking by terminal features can this kind of effect be accommodated within a bottlebrush theory. This section shows cases where non-terminal nodes need to be invoked as blockers of assimilation.

3.2.1 Coronal Transparency

First let us review some examples from Fula studied by Paradis & Prunet (1989) which indicate that vowels may spread across coronals, but not across labials and velars. Paradis & Prunet have taken such facts to argue for deep underspecification of coronals, in contrast to labials and velars.

Similar to the generalization seen above in Guere, where vowels may be fused across coronal consonants but not across consonants with other places of articulation, Fula has vowel harmony processes which apply only across coronals.

There are voice and aspect morphemes which are suffixed as empty X-slots, and derive their quality from the preceding vowel, provided the intervening consonant is coronal (52).

52. *Imperfect Paradigm*

a.	at - x	→	ata	Active voice
b.	ot - x	→	oto(o)	Middle voice
c.	et - x	→	ete(e)	Passive voice

In the pronominal system, /a/ (and sometimes /e/) assimilates to the following vowel across the coronals / / or /n/ (53).

53. *Pronouns*

a.	a - en	→	e en	*we (incl.)*
b.	a - on	→	o on	*you (pl.)*
c.	on - en	→	onon	*you (independent pl.)*

Futankoore Fula has two epenthetic vowels, /i/ and /u/, which are used to break branching onsets and codas. /i/ and /u/ can be inserted in different, but overlapping contexts. Normally, when the environment for epenthesis is met /i/ or /u/ is possible, /u/ is preferred. However, when the underlying form contains a cluster which requires insertion of *two* epenthetic vowels, as in (54), and the choice is between inserting /i/ at either side of a coronal, or inserting /i/ and /u/, the normal preference for /u/ insertion is overridden and a second /i/ is inserted. Paradis & Prunet suggest that this surprising reversal of preference is due to the fact that the two /i/s seen in the surface form are the result of just one instance of epenthesis followed by spreading to the position of the second /i/.

54. *Epenthetic vowels*

a. /ıutt -   - t/ →   ıutti it

Notice that the coronal transparency effect in Fula involves high, low, and mid vowels, and both stop and liquid coronals.

The point of these examples is to show that there are harmony processes which are blocked by just a subset of the consonants of a language, and that what distinguishes the blocking consonants from the non-blocking consonants is very unlikely to be the fact that blocking labials and velars have *all* the terminal features of vowels, while the consonants that are transparent to vowel harmony — the coronals — do not bear any of the terminal features of vowels, which is what such effects demand in a terminal feature spreading approach like Halle's. In order to capture the blocking of harmony we need to make reference to non-terminal nodes which the harmonizing vowels have in common with the labials and velars.

Different explanations have been offered for the contrast between coronal and non-coronal consonants. One popular approach is to assume that coronals are temporarily lacking a [Place] node during derivations, so that if vowel spreading is instantiated as spreading of the [Place] node, labial and velar consonants will block harmony due to violation of the Line Crossing Constraint, but coronals will not cause similar violations, as (55) shows (from Paradis & Prunet 1990).

55a.

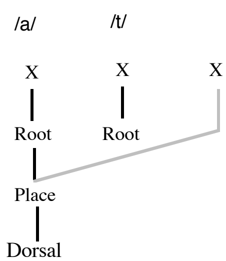

b.

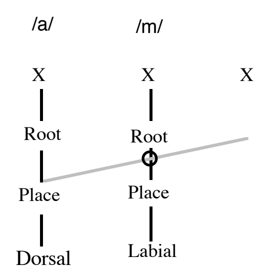

Another approach to the special status of coronals is to assume that there is a node in the feature hierarchy which exhaustively dominates the place features of labials, velars and vowels. Such a node — known as 'Peripheral' — has been suggested by Rice & Avery (1991). Given this extra structure, a process which spreads the [Peripheral] node of vowels will be full vowel harmony which will be blocked by labials and velars, but not by coronals.

There is a growing literature which raises difficulties for both of these approaches. For example, a number of people have pointed out that phenomena which treat coronals as somehow 'special', including transparency effects, tend to single out just a subset of coronals, while other coronals pattern with labial and velar consonants (cf. McCarthy & Taub 1992; Kaun 1993; Steriade 1993a). Kaun (1993) reports that a survey of coronal transparency effects shows that only /l,r,n/ are transparent to assimilation processes. Coronal obstruents and other continuants are no more transparent than labials and velars[19]. Notice, however, that these criticisms of particular analyses of the coronals do not affect the main point I am drawing from the transparency effects — namely, the need to invoke blocking by non-terminal nodes. The theories of underspecification may be inadequate, or the feature hierarchies may be incorrect, but we still expect that the appropriate characterization of the blocking and transparency effects will involve autosegmental blocking at the level of non-terminal nodes.

3.2.2 Translaryngeal Transparency

Another class of examples where vowel harmony is blocked by some of the consonants of a language but not others is the translaryngeal vowel harmony found in many languages and discussed by Steriade (1987c). In these languages vowels can assimilate in all of their oral features when they are string adjacent or separated by just /h, /. *All* oral consonants, including coronals, block assimilation. One of the most interesting cases of translaryngeal harmony discussed by Steriade is Acoma (Miller 1965), in which vowels can be glottalized or devoiced. In V V sequences in this language vowel harmony only matches the oral features of the two vowels, leaving the laryngeal features free to contrast. (56) shows V V sequences in which the vowels contrast in glottalization; (57) shows V V sequences in which the vowels contrast in voicing[20].

56. a.	ka a usiustya	*he is tied up*
b.	pu u ukaca	*come out and look at the two of them*
57. a.	ziyu uc e†e† e	*they took him*
b.	senaa a†si	*my arch*

This example shows that it would not be sufficient to claim that larygeal consonants are simply neutral to vowel spreading. The node that spreads must be [place], so that all oral consonants —which have [place] nodes — block

[19]Examples like (52) and (54) above appear to be counterexamples to Kaun's generalization that only /l,r,n/ are transparent, although their derivation is not clear. However, the point I am making here holds independent of the accuracy of Kaun's generalization.

[20]Vowel glottalization is an underlying property in Acoma. Voicing contrasts on vowels, however, are derived by rule: vowels are devoiced obligatorily after an aspirated onset (57a), and optionally before a voiceless segment (57b).

vowel spreading, and the glottalization and voicing features do not interfere with vowel spreading.[21, 22].

4. Diphthongization & Contours

4.1 The ‘diphthongization paradox’

In this section we turn to a discussion of an earlier challenge to the autosegmental theory of feature structure, raised by Hayes (1990). Hayes regards diphthongization processes as insurmountable problems for the standard interpretation of feature hierarchies. His argument runs as follows. Diphthongization processes convert sequences of 2 identical segments (long segments) into 2 non-identical segments. Autosegmental theories of feature structure represent long segments as a single feature tree which is doubly linked to the skeletal tier, as in the representation of /pp/ in (58).

[21] There are troubling instances of languages with rules which spread only across laryngeal consonants, i.e. there is a strong case for analyzing these processes as class node spreading, and yet only specific vowels may spread; this appears to require specification of a ‘dependent context’ in the rule formulation, which I have tried to avoid (cf. section 2.1 above). See Steriade (1987c) for details.

[22] An alternative would be to attempt to provide a phonetic explanation for the transparency of laryngeal consonants to vowel harmony, in order to avoid recourse to class node spreading. Stemberger (1993) discusses the feasibility of this move. Citing Keating (1985), Stemberger points out that from an articulatory point of view, glottal constrictions do not interfere with vowel-to-vowel coarticulation, but neither do labial constrictions. Therefore, articulatory facts cannot explain why laryngeal transparency should be commonly found, whereas labial transparency (to the exclusion of coronals and velars) is unknown.

On acoustic grounds, however, we might expect [p] and [h] to have different effects on vowel spreading. The labial closure of [p] causes vowel formant transitions, and during the silent closure the vowel formants of course disappear. On the other hand, [h] has low amplitude noise that is easily shaped by the vocal tract, so that the formants of adjacent vowels are easily recoverable from [h]. However, the acoustic transparency of [h] presumably does not generalize to [ ], which involves complete glottal constriction.

See Stemberger (1993) for further evidence against the perceptual account of laryngeal transparency, based on the acquisition of English phonology.

58.

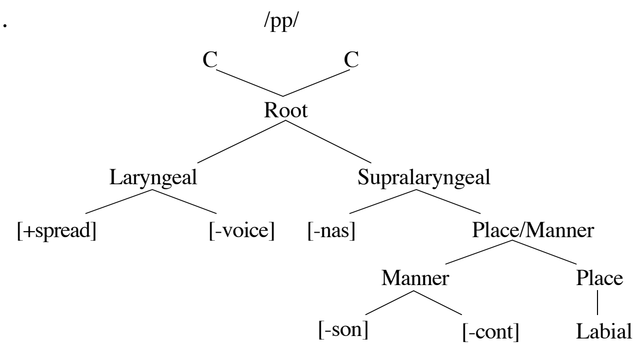

Since each feature is represented only once for the two timing slots, if any part of the feature tree is altered, it will affect both halves of the long segment. Therefore, it is predicted to be impossible for a rule to change just one half of a long segment, leaving the other half unaffected. But this is precisely what diphthongization rules do. If the characterization of diphthongization is correct, then the standard autosegmental representation of long segments is in need of revision. This is what Hayes proceeds to do. This section reviews the diphthongization facts which led Hayes to propose an alternative theory of feature structures, and suggests that the facts do not entail the fatal consequences for autosegmental hierarchies that Hayes claims they do.

> Hayes (1990):
> *Autosegmental approaches lead to the 'Diphthongization Paradox', which a Bottlebrush theory can avoid.*

Hayes was not the first person to worry about how to reconcile feature hierarchies with diphthongization rules. Clements (1985) and Steriade (1987a) had already recognized exactly the concern that Hayes raises, and had suggested modifications to the feature geometry accordingly. Both Clements and Steriade attribute the diphthongization problem to the sharing of root nodes, and so suggest analyses of diphthongization rules in which root nodes are not shared. For example, Clements treats the Icelandic preaspiration rule, which converts underlying /pp/ to [hp] (59: Thráinsson 1978), as an instance where the laryngeal and supralaryngeal tiers are separately shared by each half of /pp/. This makes it possible to delete the supralaryngeal features of the first half of the segment without affecting the laryngeal features, changing /p/ to /h/ (60).

59.	a.	kappi	[k^{h}ahpI]	*hero*
	b.	hattur	[hahtYr]	*hat*
	c.	˜akka	[θahka]	*thank*

60.

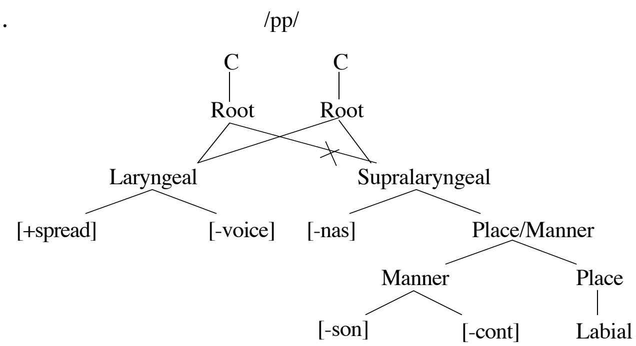

Steriade (1987a) takes a more radical position than this: she abandons altogether the notion of a single root node, and proposes that there are 3 different ‘root nodes’, each of which supports a separate subset of features and is directly associated to the prosodic tier. By separating 3 classes of features as in (61), it becomes possible to capture 3 kinds of diphthongization processes: (a) rules affecting laryngeal features, while leaving all else intact, as in Southern Paiute, where /mm/ is realized as /mm</ in certain contexts; (b) rules affecting nasality alone, as in Japanese, where /bb,dd/ can be realized as /mb,nd/; (c) rules affecting place features and nothing else, as in the case of Icelandic preaspiration.

61.

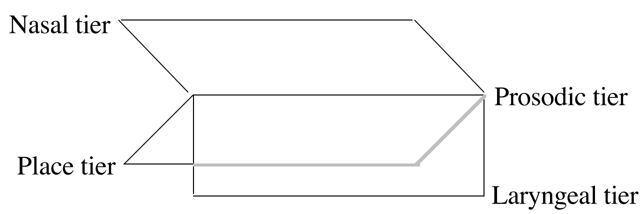

Hayes points out that Steriade’s model is only successful in reconciling diphthongization processes with a (weakened) version of autosegmental feature structure provided that diphthongization rules can in general be formulated as disrupting *all* features under one of the major class nodes and leaving all features under the other class nodes unaffected. He observes, however, that examination of a wider range of diphthongization rules shows that they can affect just individual features *within* a major class (eg. [place]), leaving all other features in that class unaffected. For example, Old French has a diphthongization process [e: o:] _ [ei ou], which affects just [±high] (Berschin *et al*. 1978).

It is of course always possible to apply Steriade’s strategy in dealing with such cases: partition those features which are and are not affected by a given diphthongization rule, so that they are separately associated to the skeletal tier. But if, as Hayes claims, *any* feature can be altered by a diphthongization rule, we arrive at a serious problem: extending Steriade’s approach to such cases leads to an abandonment of the hierarchically organized feature geometry, since every feature will have to be directly associated to the skeletal tier, and amounts to a return to the Bottlebrush model.

While Hayes denies the adequacy of the autosegmental view of feature hierarchies for diphthongization rules, he does not want to give up the successes of feature hierarchies in predicting natural classes of features. This demands a hybrid model, in which hierarchical structure is represented independently of autosegmental structure, similar to Halle's proposal.

Hayes claims that a hybrid model is motivated independently, in order to avoid confusing the dual functions of association lines in feature theory as indicating (a) simultaneity of elements on different autosegmental tiers, (b) category membership. The paradigm example of use (a) is tone association (62a), and for (b) it is syllable structure (62b) (examples from Bird (1991)).

62. a. b.

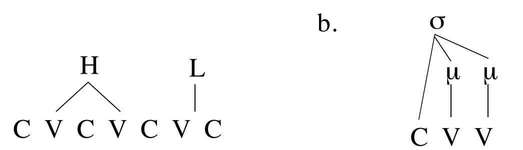

Tone: Coordination/Simultaneity

Syllables: Category Membership

Hayes proposes a version of the Bottlebrush model augmented with grouping information in the form of labelled bracketings. It is similar to the augmented Bottlebrush already seen in (8) above, which corresponds to Halle's separation of association and category membership. But it differs from (8) in that class nodes are not *purely* names for groupings of autosegmental terminal features. In Hayes' model, class nodes are still units on autosegmental tiers; they are terminal nodes, directly associated to the skeleton like the features, as in (63).

63.

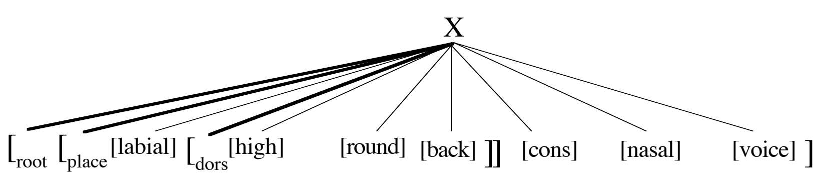

Hayes' model, unlike Halles, includes an autosegmental tier for every node in a corresponding feature hierarchy, not just the terminal nodes. Hierarchical information is only encoded in the bracketings. This means that class nodes *can* in fact be shared under this version of the Bottlebrush view, unlike in Halle's theory. Only sharing of class nodes in a model like (85) no longer entails assumption (5), repeated below.

5. *Node Sharing*
Sets of features (corresponding to subtrees) may be 'shared' by two segments by means of double linking.

The reason for this is clear from the representation in (63): class nodes do not dominate the tiers which they group: therefore, given the formal properties of the representation alone, class node sharing does not automatically entail feature sharing. Hayes claims that this is a good consequence: "inheritance of linkages from mother to daughter nodes appears to be only the normal, expected case, not an invariant law" (1990, 41). A *Percolation Convention* (64) is introduced in order to enforce the inheritance of linkages from class nodes to nodes in the same group as a default. Class nodes bear the indices of the prosodic timing slots that they are associated to, and the *Percolation Convention* ensures that a class node bears the same indices as the terminal nodes grouped with it.

64. *Percolation Convention* (Hayes 1988, 41)
When linkages are assigned to or removed from a node N[23], the assignments and deletions are automatically carried over to all nodes dominated by N.

Recall that the diphthongization paradox arises when there is a node which does not itself branch, but which is linked to multiple skeletal slots, by virtue of being linked to higher branching nodes. In Hayes' theory, since all nodes which are linked to two skeletal slots are themselves branching nodes, the paradox does not arise. The task of the *Percolation Convention* is to recreate the effects of dependencies from tier to tier which come for free in the standard Clements/Sagey model.

Since Hayes' version of the Bottlebrush theory does allow class node sharing, the arguments in sections 2 and 3 above against Halle's theory do not affect Hayes' proposal. Halle's theory is more restrictive than the Clements/Sagey model, because it excludes class node sharing. Hayes' model, on the other hand, is *less* restrictive than the Clements/Sagey model, because it accommodates all of the structural constructs of that model, plus the structures it requires to account for diphthongization rules. Therefore, different arguments are required to defend autosegmental hierarchies against Hayes' proposal.

Bracketing of features and the *Percolation Convention* are devices which recreate automatic properties of the autosegmental hierarchy model. This complicated way of replicating the effects of autosegmental hierarchies is presumably not desirable unless necessary. In the sections that follow I reconsider the arguments which lead Hayes to abandon autosegmental hierarchies in favour of his enriched Bottlebrush model.

[23]'Linkages' in this definition are indices corresponding to elements on the CV tier. Hayes does not actually use representations like (63) in most of his discussion - he adopts a different notation for reasons of graphical tractability in dealing with sequences of segments. Details of this alternative notation are not important to the discussion here, and the representation in (63) illustrates most clearly what Hayes considers the hierarchical aspect of feature structures to consist of.

4.2 What does diphthongization affect?

If we accept Hayes' assumption that (i) diphthongization rules take doubly linked feature structures as inputs, and (ii) that diphthongization rules can affect any part of a feature structure, then his arguments for a return to the Bottlebrush theory are quite compelling. However, it is useful to reconsider the evidence leading to this conclusion.

First, in arguing for associating all tiers directly to the prosodic tier, Hayes says of Steriade's approach: "...her theory does not go far enough, because diphthongization is not limited to the major class nodes [...] Many diphthongization rules affect just a single feature, deeply embedded in the tree." (Hayes 1990, 38) This implies that diphthongization can affect *any* node in a feature structure, just as harmony processes are assumed to do (cf. McCarthy 1988).

However, a look at the variety of attested diphthongization rules suggests that Hayes' point should be made differently. There are certainly plenty of cases of diphthongization rules that affect individual features; what is hard to find is diphthongization rules which *do not* affect just individual features. Of the three examples cited from Steriade (1987a), two are expressible in terms of changes at the level of terminal features — /nn/ _ /nn†/ in Southern Paiute affects just the laryngeal feature [±voice], and /bb,dd/ _ /mb,nd/ in Japanese affects just the feature [±nasal]. The third case, Icelandic preaspiration, is discussed at length in the next section. Table 1 below shows that a wide range of cases discussed by Hayes all affect individual terminal features.

Table 1: Diphthongization rules as operations on terminal tiers

Middle English (beginning of Great Vowel Shift	[i: u:] _ [ei ou]	Add [-high] left
Old High German (Penzl 1969)	[e: o:] _ [ea oa] _ [ia ua] _ [ie uo]	Add [+low] right; add [+high] left; delete [+low] right
New High German (Chambers & Wilkie 1970)	[i: u:] _ [ai au]	Add [+low] left (repair inserts [-high] left
Icelandic (Garnes 1974)	[e: o: fl:] _ [e« o¿ fl˘]	Add [-ATR] right
Faroese (Rischel 1968)	[u: o: i: y:] _ [yu ou ui ui]	
Central Swedish (Elert 1981)	[i: y: ,: u:] _ [ij y	Add [+cons] right
Scanian Swedish (Bruce 1970)		
Mälaren Region Swedish (Elert 1981)	[V:] _ [V\]	see below
Old French (Berschin et al. 1978)	[e: o:] _ [ei ou]	Add [+high] right
Quebec French (Dumas 1981)		
Slovak (Kenstowicz & Rubach 1987)	[e: o: :] _ [ie uo ia] in derived environments	Add [+high] left. Repair fronts [æ] to [a]
Czech Common Language (Kucˇera 1961)	[e: o:] _ [ei ou]	Add [+high] right
Old Prussian (Schmalstieg 1964)	[e: o:] _ [ei ou]	Add [+high] right
Eastern Finnish (Kiparsky 1968)	[e: fl: o: : a:] _ [ie yfl uo oa]	Add 1 degree of height left
Southern Lappish (McCawley 1973)	[i: y: u: e: o:] _ [ie yfl uo oa]	Remove 1 degree of height right
Cakchiquel (Campbell 1971)	[e: o:] _ [ie uo]	Add [+high] left

As can be seen from a glance at the table, the vast majority of cases are straightforwardly analyzable as operations adding or deleting a single terminal node, either in terms of binary features, or degrees of height. Only in the case of Mälaren Region Swedish does an account referring to operations on class nodes suggest itself; since the diphthongization involves complete neutralization of the second half of the long vowel, we might analyze the rule as delinking of the root node from the second timing slot, followed by insertion of the default vowel /\/.

4.2.1 *Icelandic Preaspiration: A Counterexample?*

The case which seems to best motivate the claim that diphthongization can affect major class nodes is the Icelandic preaspiration rule introduced above. Thráinsson (1978), Clements (1985) and Hayes (1990) analyze this process as delinking of the [supralaryngeal] class node. Since I am predicting that diphthongization rules in Hayes' sense only apply at the level of terminal features, Icelandic preaspiration presents a serious difficulty for my generalization. This section reexamines the arguments for the diphthongization analysis.

The Icelandic preaspiration rule (Thráinsson 1978) converts underlying /pp,tt,kk/ sequences to /hp,ht,hk/. If we regard this as a rule which affects the features [±sonorant], [±continuant] and [coronal] in the first half of the consonant cluster, we have a strong case for a diphthongization rule affecting a major class node, since there seems to be no motivation for regarding preaspiration as three separate operations, affecting each of these features individually.

The force of Hayes' criticism of autosegmental class node sharing depends on the premise that diphthongization rules change structure which is shared between two segments, but that this structural change only affects one of the two segments. In this section I argue that Icelandic preaspiration does not conform to this pattern. I suggest that preaspiration is a reordering effect motivated by contrast licensing, rather than an effect of structural change.

Two facts about Icelandic phonology are central to my analysis here. First, the contrast which is generally represented by /p,t,k/ versus /b,d,g/ is not one of voicing, but one of aspiration: /p,t,k/ are voiceless aspirated stops, whereas /b,d,g/ are voiceless unaspirated stops. Following Thráinsson (1978) I assume that /p,t,k/ are distinguished by the feature [+spread glottis]. The aspirated stops involve an active spreading of the vocal folds, leading to a wide open glottis. The effect of the glottal spreading, according to Kim (1970) is that following the release of the oral closure of the stop the vocal folds cannot immediately assume a position normal for voicing the succeeding vowel. Hence, the *perceptually salient* effect of the glottal gesture is a period of aspiration between oral release and the onset of voicing.

The importance of this point is that the feature [+spread glottis] which distinguishes /p,t,k/ from /b,d,g/ is not expressed simultaneously with the other features of the stops, such as place and the supralaryngeal manner features.

Second, preaspirated stops in Icelandic do not only arise from underlying /pp,tt,kk/, as in (59) above. They also arise from underlying /p,t,k/ when followed by /l,n/. (65) shows examples without alternations, and (66) shows synchronic alternations between [p^h] and [hp] when the genitive plural suffix /-na/ is added to stems ending in /p/.

65.	a.	epli	[«hplI]	*apple*
	b.	ætla	[aihtla]	*intend*
	c.	opna	[¿hpna]	*open*
	d.	vakna	[vahkna]	*wake up*

66.			Nom. sg.		Gen. pl.	
	a.	pípa	[$p^hi:p^ha$]	pípna	[p^hihpna]	*pipe*
	b.	gata	[$ka:t^ha$]	gatna	[kahtna]	*street*
	c.	kaka	[$k^ha:k^ha$]	kakna	[k^hahkna]	*cake*

Ideally, we would like to provide an analysis of preaspiration which generalizes to the two environments in which it occurs.

Once we take into account the fact that aspirated stops in Icelandic can be split into two subparts (67a), it becomes clear that preaspirated stops are not just pairs of two aspirated stops (67b) with some structural change applied to the first half (67d), contrary to Thráinsson's characterization of them. Even if we group the closure and release portions of the two segments, as an autosegmental representation like (58), predicting (67c), there is no way of deriving the form of a preaspirated stop (67e) by just changing the first half of the geminate.

67.	a.	Aspirated stop	p^h	Closure+Release
	b.	Geminate aspirated stop	$*p^hp^h$	Closure+Release+Closure+Release
	c.	Geminate aspirated stop	$pp^{h(h)}$	Closure+Closure+Release
	d.	Preaspirated stop	$*hp^h$	Release+Closure+Release
	e.	Preaspirated stop	hp	Release+Closure

The first half of a preaspirated stop expresses the features of the release portion of a single aspirated stop, and the second half expresses only the closure features of the aspirated stop. Crucially, *pre*aspiration involves loss of *post*aspiration. If preaspiration is the diphthongization rule described by Thráinsson, Clements and Hayes, we incorrectly expect the output to be the pre- and postaspirated stop in (67c).

Looking at the descriptions of gestures in (67a&e) above, the most obvious analysis of preaspirated stops is as a realignment of the two components of aspirated stops, which assumes 'nothing added, nothing taken away'. If we assume that the only ordering imposed by autosegmental representations is ordering on individual autosegmental tiers (plus the effects of the Line Crossing Constraint), then structures like (58) do not impose any ordering on the feature [+spread glottis] with respect to the supralaryngeal features.

If we assume that geminate stops have doubly linked root nodes, such that each feature is realized only once across the two timing slots, preaspiration arising from /pp/ and /pn,pl/ will manipulate identical structures.

What the two inputs to the preaspiration rule share is that in each case there is a coda segment which is associated with an aspirated stop, but which lacks a C-V transition in which to realize its [+spread glottis] feature. In this situation the [+spread glottis] feature is realigned with respect to the closure features of the coda stop so that it precedes the closure[24].

---

[24]Kingston (1990) argues that the glottal articulation in Icelandic preaspirated stops does not correlate with the following oral closure. His argument is based on the finding that changes in the duration of glottal abduction in fast speech covary with changes in the duration of the preceding vowel rather than changes in the following closure.

This does not affect my claim that preaspiration is reordering of the expression of oral features and [+spread glottis] for a couple of reasons. First, Kingston also reports a lack of correlation between changes in the duration of closure and glottal abduction for postaspirated stops in Icelandic, making it hard to draw conclusions about the special status of glottal/oral coordination in preaspirated stops. Second, my account of reordering in preaspirated stops is based on perceptual factors, and so it is the duration of aspiration rather than the duration of glottal abduction that we should regard as important from a perceptual perspective (glottal

This reordering of closure and release features in environments where their canonical ordering is impossible can be viewed as an analogue of Steriade's *Positional Neutralization* (Trubetzkoy's 'Strukturbedingte Aufhebung'). The phonology pays attention to whether its output will be easily recoverable by the listener; it has two possible responses to the problem of loss of perceptual salience of consonant features in certain positions. If the relevant feature is expressed simultaneously with other features, the contrast it expresses is neutralized. If, on the other hand, the relevant feature is ordered with respect to other features, and can be adequately expressed by reordering, the contrast will be rescued in this way rather than neutralized.

The view of preaspiration as structure-reordering rather than structure changing is by no means new (eg. Haugen 1958, 72). A number of arguments have been raised against it, though, some of which I address here.

a. Length of aspiration.

First there is the issue of the length of preaspiration. As reflected in the fact that preaspiration is generally represented as /h/, whereas postaspiration is represented as [^{h}], preaspiration generally lasts a good deal longer than postaspiration (Petursson 1972, Garnes 1974). Based on the length contrast Thráinsson assumes that "this implies the giving of a segment status to preaspiration and not to postaspiration, and a rejection of the view sometimes advanced that preaspiration is to be considered a component or a phonetic feature of the succeeding stop" (Thráinsson 1978, 5).

I find this dismissal rather hasty, given that there are plausible phonetic reasons for this phonetic contrast. Consider first the *articulatory* status of pre- and postaspiration, assuming that they are the expression of the feature [+spread glottis]. In a postaspirated stop the vocal folds can be spread during the closure portion of the segment and remain wide open throughout the duration of that gesture. The measured period of aspiration begins at the release of the stop and lasts until the vocal folds can attain a position appropriate for voicing. In a preaspirated stop the time required to spread open the vocal folds is no longer contained within the closure period, so it must be calculated in with the length of aspiration. In this case too the musculature involved in spreading the vocal folds has to oppose the force generated by vibration of the vocal folds, and the intensity of aspiration is not enhanced by the pressure build-up during closure which postaspirated stops benefit from. The effect of this is surely that a more exaggerated articulatory gesture is required for preaspiration than postaspiration, which is likely to have an effect on the length of the gesture.

From a perceptual perspective there are also good reasons to expect a length difference between pre- and postaspiration. Postaspiration has an extremely well defined onset, at the release of the stop closure, and a large body of research indicates that adults, young infants[25] and non-human species[26] are remarkably sensitive to the time between stop release and onset of voicing, accurate to within a couple of milliseconds. So postaspiration is easily anticipated following a period of silence; it is easily perceived following that silence, and there appear to be highly specific innate mechanisms for analyzing it. On the other hand, preaspiration cannot be anticipated from the vowel which precedes it, it does not have such a well-defined onset as as the stop release (burst of frication) which marks the start of postaspiration, there is no evidence for specific perceptual mechanisms for analyzing preaspiration; and since it follows immediately from a period of voicing, it will stand out less against its acoustic 'background'. All of which is supposed to argue that there are plenty of reasons why expressing [+spread glottis] before rather than after the stop closure should make an enormous difference to the length of aspiration. Therefore, the fact that aspiration is longer when it precedes the stop closure than when it follows it follows the closure does not automatically imply that preaspiration has segmental status.

b. Geminate Input.

A second argument raised against the 'inversion' analysis of preaspiration derives from the observation that aspirated stops do not become preaspirated in just any contexts. It has been claimed that cross-linguistically preaspiration only applies to underlying geminates (Thráinsson 1978; Kingston 1985; Keating 1990). Given a geminate input, then, it is attractive to view preaspiration as a rule which simply deletes the supralaryngeal features of the first half of the geminate, leaving a segmental /h/.

An obvious challenge to the 'geminates only' generalization about preaspiration is the underlying stop+liquid sequences in Icelandic which also lead to preaspiration. Thráinsson gives arguments that the stops in these sequences undergo lengthening prior to preaspiration, so that the geminates only analysis can be maintained. However, the evidence for lengthening of /pn/ to /ppn/ before preaspiration is inconclusive; plus, the motivation I gave for preaspiration as 'inversion' explains why preaspiration does not apply to just any stop — it just applies to unreleased coda stops — and why the rule does what it does — it orders the [+spread glottis] feature so that it can be expressed — without recourse to the geminates only condition.

---

abduction overlaps with closure). Kingston reports that duration of preaspiration and duration of closure *do* covary.

[25]Eimas et al. (1971) demonstrated that English-learning infants aged 1 to 4 months show far better discrimination along a voice onset time continuum when the pairs of stimuli straddle the /ba/ - /pa/ category boundary than when they both fall within the range of one category. For reviews see Aslin 1987; Kuhl 1987).

[26]Kuhl and Miller (1975) showed that chinchillas exhibit category-like identification of aspirated and unaspirated stops. See Sinex and McDonald (1989) for evidence that this categorial discrimination of stops in chinchillas may be the result of very low level processing in the auditory nerve.

Thráinsson presents a number of phonetic measurements in support of his claim that stops preceding /l,n/ undergo lengthening before undergoing the preaspiration rule. He first shows that preaspirated stops preceding liquids have the same length, whether they correspond to underlying geminate+r or singly or doubly spelled stop+n sequences (68-69).

68.

Preasp.+stop / __n		Preasp.+stop / __r		Short stop /__r	
vatna	213	settra	205	setra	120
vakna	197	dekkra	205	dekra	95
Mean:	205		205		108

69.

	Underl. long stop				Underl. short stop		
	[h]	C	[h]+C		[h]	C	[h]+C
heppni	74	120	194	vopna	104	124	228
hittni	112	133	245	vatna	90	123	213
˜ykkna	100	117	217	vakna	117	80	197
Mean:	95	123	219		104	109	213

The fact that the surface forms of preaspirated stops have the same length does not have to be attributed to a common geminate substrate. We have already seen reasons why preaspirated stops are expected to be much longer than postaspirated stops.

Thráinsson uses a second set of measurements to argue for a general stop lengthening rule before /l,n/, regardless of whether the stop is aspirated. He compares the length of the voiceless unaspirated /b,d,g/ before /l,n/ with the length of *underlyingly* geminate /bb,dd,gg/. This test is more interesting for our purposes, since my account of preaspiration makes rather different predictions from Thráinsson's here. If preaspiration in stop+liquid sequences is simply driven by the need to express [+spread glottis] in the coda, as I claim here, there is no reason to expect [-spread glottis] stops to be affected in the same environment. Thráinsson, on the other hand, predicts that stop lengthening before /l,n/ applies regardless of whether the stop is aspirated.

Unfortunately the results of these measurements, shown in (70) are hard to interpret, given the difficulty of finding minimal pairs of short and geminate unaspirated stops in the relevant contexts in Icelandic. The comparison manipulates two variables at the same time. The length of underlyingly short /b,d,g/ before /l,n/ is compared to the length of underlyingly geminate /bb,dd,gg/ before /r,v/ (70). The underlyingly short stops turned out to be *longer* than the underlying geminates, although the data are not clear. This is not expected under any account.

70.

Underl. long stop / __ r,v		Underl. short stop /__l,n		
		tefla	[tebla]	184
saddra	150	neinna	[neidna]	203
sneggra	125	ragna		124
glöggva	115	mygla		117
Mean	130			157

So it appears that Thráinsson's data do not yet distinguish between his gemination-followed-by-diphthongization account of stop+liquid sequences and my inversion account. But as we have seen, a more extensive study of the behaviour of /b,d,g/ preceding /l,n/ could provide strong evidence against my analysis. In the absence of such data, however, the inversion analysis is attractive, because it explains *why* preaspiration applies to short aspirated stops only when they are followed by liquids with an alveolar constriction. It also does away with the need for ordering the preaspiration rule with respect to other rules. The diphthongization analysis requires that preaspiration follow pre-sonorant lengthening; but preaspiration must precede the word-final devoicing rule, which turns /bb,dd,gg/ into /pp,dd,gg/, because these underlying geminates surface as geminates rather than as preaspirated stops. Under my account, the devoiced geminate stops are not candidates for preaspiration, since they have no [+spread glottis] feature to realize.

71.

	/hattYr/	/kr¿bb/	/«pli/
Lengthening /__l,n	-	-	«ppli
Preaspiration	hahtYr	-	«hpli
Final Devoicing	-	kr¿pp	-

In sum, Icelandic preaspiration appears no to require an analysis as a diphthongization rule affecting just the first half of a long segment. Therefore, it is quite consistent with the generalization that diphthongization rules affect only terminal features.

4.3 Diphthongs as Contour Segments

The generalization that true diphthongization rules only affect terminal features is surprising in the light of Hayes' theory[27]: Hayes claims that diphthongization rules can apply at any point in a feature hierarchy, and the representations he proposes are designed to accommodate this expectation. Meanwhile, there is a possible analysis of diphthongization which both predicts our generalization about possible diphthongizations and retains standard

[27] What predictions about diphthongization are made by Halle's version of the bottlebrush theory is less easy to determine, since it's not clear what use he intends to make of the feature trees in his representations. But given the reasonable inference that autosegmental operations like delinking and structure insertion are only permitted at the level of terminal features, as Halle assumes for spreading and feature sharing, Halle makes identical predictions about diphthongization to the autosegmental approach I have been advocating here.

assumptions about autosegmental feature hierarchies and class node sharing. This is the analysis of diphthongs as contour segments, based on Sagey's (1986) account of contour segments, as in the following repreesentation for the diphthong /ei/, taken from Hayes' paper.

72.

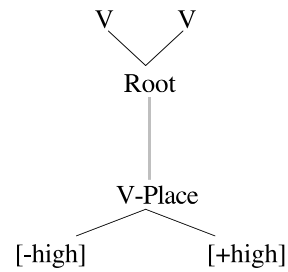

Under this view, a diphthongization rule simply appends a terminal node to a feature structure. For example, the rule /ee → ei/ appends [+high] to the right of a doubly linked [-high]. This seems to be a simple analysis of the phenomena, and is very similar to a standard view of affricates as contour segments, but there are a number of possible arguments against it, which I will try to deal with here.

4.3.1 Length

Hayes anticipates the contour segment analysis of diphthongs, but rejects it because it fails to account for the generalization that diphthongization rules only apply to long vowels, and that diphthongs in general, whether or not they are generated by a synchronic rule, almost always occupy two prosodic positions. Hayes correctly points out that nothing in autosegmental representations of contour segments predicts this, and that a constraint on feature structures which imposed this restriction would need to be a rather unnatural one — it would have to say something like 'terminal nodes may bear two values iff the root node is doubly linked.'

First, it should be noticed that Hayes' *representations* make the same predictions about the length of diphthongs as the standard hierarchical autosegmental representations, i.e. none. There is no principled way — either in Hayes' account or in the contour segment analysis I am proposing — to differentiate affricates, which quite often occupy a single prosodic position, from diphthongs, which are generally long. This is probably a good result, however, given that Hayes' more robust generalization is that active *rules* of diphthongization apply to long segments only. Since the rules themselves have to encode whether they apply to short or long segments, they will carry the burden of capturing this generalization, and not representational constraints *per se*.

What the phonological representations *are* more responsible for is accounting for any absolute length restrictions on diphthongs, independent of whether they are the output of synchronic diphthongization rules. Here, however, we find a number of languages with short diphthongs — examples are Ancient Greek (Steriade 1988), Chevak Yupik (Woodbury 1981), Icelandic (Anderson 1984) and Faroese (Rischel 1968) — which ought to be sufficient motivation for allowing representations to encode short diphthongs.

There is a clear distinction between trying to design universal feature representations so that it is impossible to represent short diphthongs, and observing that short diphthongs are rather rare. In this second regard, Hayes' observations about the length tendencies of diphthongs are well taken, but we should probably look somewhere other than our representational devices for an explanation of this fact, and of the contrast with affricate consonants, which don't show the same tendency to occupy two prosodic positions.

An explanation for the length of diphthong vowels compared to affricates might lie in the inherent articulatory properties of the features they realize. Perceptually, vowels are predominantly differentiated by the formant values at the centre of the vowel, which require length in order to be identified, whereas the most perceptually informative portions of obstruents are typically the transitions at their periphery, and the closure itself is unimportant. Articulatorily, repositioning of the tongue from one vowel target to another in a diphthong involves much more movement than the transition from one degree of aperture to another in an affricate. These crude facts alone would lead us to expect diphthongs to be longer than affricates.

So, Hayes' own objection to the contour segment analysis of diphthongs can be accommodated. As we have already seen, the contour segment analysis makes different predictions from Hayes about the variety of possible diphthongizations, predictions which appear to be correct.

4.3.2 What are Contour Segments?

Another potential problem for the claim that diphthongs are contour segments is posed by arguments that contour representations like (100) should be abandoned. In this section I try to defend my proposal against such objections.

In a series of papers Steriade (1992a,b) has argued that contour segment representations like (72) fall well short of explaining the variety and distribution of contour segments. Building on the generalization that only plosives can have distinctive intrasegmental contours, Steriade (1992b) proposes that plosives have bipositional representations as in (73) — in which each position corresponds to a different degree of aperture — and that various constraints on possible contour segments should follow from these representations. The three degrees of aperture which Steriade represents are A_0 (closure), A_f (turbulent airflow: fricative), A_{max} (release/approximant). Crucially for Steriade, only plosives have bipositional structure.

73. a. [nas]
A_0 A_{max}
nasal stop

b. [nas]
A_0 A_{max}
prenasal stop

c. [spread]
A_0 A_{max}
postaspirated stop

d. [spread]
A_0 A_f
preaspirated affricate

The generalizations which these representations are intended to capture are the following:

74.i. Only plosives show distinctive contours; vowels, approximants and fricatives do not.
ii. Contours tend to be lost in environments where plosives are unreleased.
iii. "Distinctive intrasegmental contours never exceed two articulatory phases." (Steriade 1992b, 3).

Generalization (i), due to Maddieson (1984), is captured by the assumption that only plosives are bipositional. If an unreleased plosive is missing one of its positions, then generalization (ii) follows. If phases of contours are restricted to one per subsegmental position, and plosives are limited to two positions, generalization (iii) follows.

I have very little disagreement with the generalizations in (74). However, I am not convinced that they warrant giving up Sagey's representation of contours in (72) in favour of Steriade's multi-positional representations in (73).

Note that generalization (74i) makes a very precise claim. What is special about plosives is not that they allow contours. If we call a segment in which different features are realized *separately* and in a *fixed* order a contour segment, fricatives and approximants do allow contours. For example, the Mazateco dialect Mazatlán de Flores (Kirk 1966, 14-25, Steriade 1992b) has preaspirated continuants and postaspirated fricatives (75c-d). What is special about plosives, according to Steriade's generalization (74i) is that the ordering of the different phases of the contour can be phonologically distinctive, as exemplified by the pre- and postaspirated plosives of Mazatlán de Flores in (75a-b).

75.a. *oral plosive + h*
th, tsh, t§h, kh, k^wh

c. *continuant + h*
sh, §h

b. *h + oral plosive*
ht, hts, ht§, hk, hk^w

d. *h + continuant*
hy, hw, hwy [h·]

The distinctive ordering of plosive contours is straightforwardly representable using either Sagey's format, as plus and minus values of a feature associated to the same node (72)[28], or in Steriade's format, as a privative feature associated to one of the two positions of a plosive (73c). The empirical difference is that Sagey's representations can encode distinctive fricative or approximant contours just as easily as plosive contours, whereas Steriade's representations cannot, because only plosives are assumed to have two positions. Steriade suggests the representations in (76) for aspirated fricatives and glides.

76.

spread	spread
|	|
A_f	A_{max}
aspirated fricative	aspirated glide

Although the ordering of aspiration with respect to other features is not distinctive for fricatives and approximants, the strict ordering must be enforced somehow, and Steriade's phonological representations do not encode the ordering (in addition, the ordering is different for aspirated glides and aspirated fricatives in Mazatlán de Flores). This task is presumably the task of the phonetics. But if the phonetics rules out the impossible orderings of approximant and fricative contours, the phonological representations do not need to impose the same restriction. Thus the criticism that Sagey's representations overgenerate non-plosive contours might not be fairly levelled against the phonological representations[29].

The second generalization, that plosive contours generally disappear in environments where plosives are unreleased (74ii), i.e. contours are lost in codas, could presumably be expressed as either a phonological deleting an A_f or an A_{max} slot, or as a phonetic consequence of the loss of release. What follows are a couple of reasons for choosing between these two alternatives, which appear to favour a phonetic account over Steriade's phonological account.

a. Phonetic and Phonological Neutralization

Phonetically, we have already seen that contours are not restricted to plosives. Phonetically, non-plosive consonants are also released. For example, Mazatlán aspirated fricatives express [+spread glottis] at the release of the

---

[28] Since I am assuming that feature tiers have varied content (see section 1.3), plus and minus values of the same feature occupy the same tier, which means that their ordering is fixed by the phonology. Unfortunately this 2-dimensional medium makes it difficult to represent the contrast between sister nodes on *different* tiers, which are not inherently ordered, and sister nodes on the *same* tier, which are inherently ordered.

[Note that my interpretation of (72) differs from that of Lombardi (1990), who assumes that plus and minus values occupy different tiers, and are therefore not ordered by the phonology.]

[29] While I am claiming that generalization (i) will have to be captured in *at least* the phonetic component whether we adopt Sagey's or Steriade's representation of contours, this does not amount to an explanation of *why* generalization (i) holds.

fricative. Therefore, both plosives and fricatives can be released or unreleased, and the articulators must be given instructions about which choice is appropriate. In Steriade's system there is an asymmetry between plosives and fricatives in regard to where these instructions are issued. For plosives, the phonology decides whether there is release or not, and the phonetics interprets the single A_0 position as a single phonetic position. For fricative contours, on the other hand, the phonology outputs a single A_f position, and the phonetics chooses whether to assign to this one or two phonetic positions.

There may be good reason to introduce this asymmetry between plosives and non-plosive consonants, but it needs to be specified. The criterion of distinctive ordering, which Steriade uses to determine which contours are encoded phonologically and which are not, cannot distinguish between plosives and fricatives here, since loss of aspiration from a plosive or a fricative has exactly the same effect of neutralizing a [±spread glottis] contrast.

b. Plosive-specific Neutralization.

Steriade's position deletion analysis of neutralization of contours in codas makes the following simple prediction. It should be possible for a language to neutralize plosive contours by failing to release them in coda position, but fail to neutralize non-distinctive fricative contours in the same positions. The most common scenario is that either all contours, contrastive or not, are lost in coda position (eg. Mazatlán), or all are retained (eg. Tzutujil, Dayley 1985). I am unaware of any examples of the intermediate situation that Steriade predicts.

c. What is deleted?

Deletion of the 'release' slot in Steriade's representations and honetic loss of release do not necessarily yield the same output. Consider the case of the prenasalized stop /^{n}d/. In Steriade's representation, the A_0 position of the stop is linked to the feature [nasal], whereas the A_{max} position is not linked to [nasal]. If these representations are taken as literal articulatory instructions, movement of the velum will be precisely synchronized with the release of closure. In fact, the velum must change position *during* the closure in order for perception of a prenasalized /^{n}d/ rather than just /n/[30]. The same is true for postnasalized stops like /d^n/, just with the opposite velar movement.

Importantly, although the velum must change position during the closure, this change is extremely hard to perceive if the stop is not released. Therefore, careful phonetic analysis could distinguish between the following two possible accounts of unreleased nasal contours.

i. Loss of release position A_{max} (phonological). The A_{max} position in Steriade's representation of a nasal contour is deleted, leaving prenasalized stops identical to unreleased /n/ (77a) and postnasalized stops identical to unreleased /d/ (77b).

[30]Burton, Blumstein & Stevens (1992) report that velopharyngeal closure precedes oral release in /^{n}d/ by around 20ms in Moru.

77. a.

$$\begin{array}{cc} [\text{nas}] & /\text{d}/ \\ | & | \\ A_0 & A_{max} \end{array} \longrightarrow \begin{array}{c} [\text{nas}] \\ | \\ A_0 \end{array}$$

prenasalized stop

b.

$$\begin{array}{cc} /\text{d}/ & [\text{nas}] \\ | & | \\ A_0 & A_{max} \end{array} \longrightarrow \begin{array}{c} /\text{d}/ \\ | \\ A_0 \end{array}$$

postnasalized stop

ii. Loss of release (phonetic). The nasal contour is not released. However, careful measurement reveals that the change in velar position takes place as normal, only it is imperceptible.

This second state of affairs would be analogous to the finding that neutralization of dental/retroflex contrasts in initial position in some languages is perceptual rather than articulatory neutralization (Anderson 1993 on Tiwi; Spajic 1993 on Toda; Steriade 1994). In initial position, the tip of the tongue goes through the same movements for retroflexes as in other contexts, but in the absence of a V-C transition the difference between retroflexes and dentals is imperceptible.

Dental-retroflex neutralization and nasal contour neutralization share the property that they could be explained either in terms of loss of perceptual contrast or in terms of phonological simplification. The learner is faced with a difficult problem if UG leaves such ambiguous situations truly ambiguous: since the learner lacks the phonetician's tools, she has no way of choosing between the phonological or perceptual neutralization account of what she hears. One response to this would be to free the learner from making the choice and assume that neutralizations are assumed to be like the Tiwi/Toda retroflex neutralization case, and involve neutralization only at the perceptual level, unless there is positive evidence for phonological simplification. If UG is helpful in this way, and never requires the learner to choose between the two analyses of neutralization in the absence of evidence, then there is no need for representations allowing conversion of prenasals to nasals, like those in (77a).

In sum, facts about neutralization in unreleased contours are potentially highly relevant for choosing between Steriade's analysis of contours and Sagey's model. I have suggested reasons for preferring a phonetic analysis. I have also pointed to potential evidence which would favour Steriade's analysis of contour neutralization as loss of a phonological position; however, more evidence is needed in order for neutralization facts to offer compelling support for Steriade's bipositional analysis.

Steriade's third generalization — that "distinctive intrasegmental contours never exceed two articulatory phases" (Steriade 1992b) — again relies on restricting attention to distinctively ordered contours. Assuming that prenasalized stops have two articulatory phases, and that postglottalized stops do too, the prenasalized postglottalized stops like /^{n}t / which function as single segments in many Mazateco dialects (Steriade 1992a,b) surely have three articulatory phases. However, the existence of (non-distinctive) ternary contours favours neither Sagey's nor Steriade's approach, so we need not dwell upon it here.

To place this discussion in the context of the main argument of the paper. Hayes claims that diphthongization poses an insurmountable problem for autosegmental theories of feature hierarchies. I have argued that Hayes' criticism of the autosegmental theories can be avoided by assuming that diphthongs are contour segments, borrowing Sagey's representation of contours (100). In order to do this, however, I have had to defend these representations against Steriade's (1992b) arguments that such representations should be abandoned. This section has tried to show that more evidence is needed before we abandon Sagey's approach to contour segments, and there are still good reasons to assume her representations.

4.3.3 Underspecification[31]

If the contour segment analysis of diphthongs can survive the general objections in 4.3.1 and 4.3.2, it has to face a more specific problem raised by underspecification theory. We have been assuming that diphthongization rules append a feature [-αF] to a segment which already bears [αF], resulting in a contour segment. Now imagine that we have independent reason to believe that [αF] is not specified underlyingly, and is only filled in as a default rule at the end of the derivation (78a). If a diphthongization rule appends [-αF] to an underspecified and doubly linked segment, this will block the insertion of [αF] by default rule, with the result that the diphthongization rule is incorrectly predicted to affect both halves of the long segment it applies to (78b).

78.
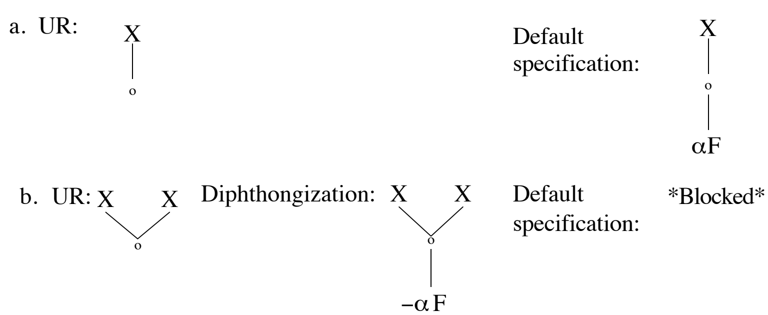

Steriade (1987a) discusses a diphthongization rule in Southern Paiute which gives rise to precisely this concern. In various environments long vowels or sonorants are partially devoiced (79)[32].

---

[31]Thanks to Donca Steriade for pointing out this difficulty to me.
[32]These devoicing facts are part of a far more complex body of devoicing data in Southern Paiute. See Sapir (1930).

79a.	/qaanna-/	*singing*	_	qaann<a<cuvaippi'a *still-his-own-singing said*
b.	/qa'aa-/	*to begin to sing*	_	qa'aa<
c.	/morii-/	*bean*	_	mor ,

If the the devoicing rule supplies the feature [-voice] to doubly linked segments in (79), and if voicing is not underlyingly specified for vowels or sonorants, we are led into the problem in (78b) — the devoicing operation is predicted to devoice both halves of the long segment, contrary to fact.

The problem is striking in the case of Southern Paiute, because do not expect to have to specify voicing distinctively for vowels. But the problem is one which could arise generally, depending on what assumptions are made about the underlying specification of features. If only *marked* values of features are included in underlying representations, i.e. radical underspecification (Archangeli 1984, 1988) then any diphthongization which appends a marked feature value to an underlying long segment will run into the same difficulty seen in Southern Paiute.

One way to accomodate examples like Southern Paiute devoicing is to assume that diphthongization rules do not always precede the specification of default values. In this way [+voice] is present on vowels before [-voice] is appended.

5. Conclusions

I have defended the view that feature hierarchies are hierarchies of autosegmental tiers. To the extent that this is successful, phonological representations are more homogeneous, given the well-established autosegmental properties of the tonal tier.

In order to maintain this position, I am forced to concede that there are some generalizations that cannot be captured by the theory of phonological representations. Provided that explanations for these generalizations can be given by theories of phonological transformations or articulation and perception, I consider this a positive result.

In an influential 1988 review article, McCarthy observes that a leading assumption of a much phonological research since the late 1970s has been that "primary emphasis should be placed on studying phonological representations rather than rules [...] if the representations are right, then the rules will follow." (1988, 84). In sections 2 and 3 I have tried to show that this heuristic of assuming that representations rather than operations or transformations are to blame can appear to force radical changes to the theory of representations, where relatively benign modification of the theory of operations is at least as effective.

McCarthy mentions another heuristic which has guided a large amount of research: "The goal of phonology is the construction of a theory in which cross-linguistically common and well-established processes emerge from very simple combinations of the descriptive parameters of the model" (1988, 84); "we should be able to freely combine the predicates of our theory of representations and our theory of operations and constraints and, in each case, come up with some real rule that languages have" (1988, 90). In other words, generative phonology should generate a class of languages which fits very snugly to the class of attested languages.

While this second idea is a useful heuristic in the construction of theories, the heuristic should not be confused with an expectation about the way the world is. Attested phonologies reflect the interaction of phonology with a number of other systems. First, it would be surprising if all phonologies generated by UG are *learnable*; second, what phonology generates has to be *usable* by articulatory and perceptual phonetic systems. Both of these interactions, which have been discussed at various points in this paper, should be taken into account when we find that our phonological theory appears to 'overgenerate'.

References

Allen, W. 1962. *Sandhi*. Mouton: The Hague.
Anderson, S. 1984. A metrical reinterpretation of some traditional claims about quantity and stress. In: M. Aronoff & R. Oehrle (eds), *Language Sound Structure*. Cambridge, MA: MIT Press.
Anderson, V. 1993. *An acoustic investigation of Tiwi consonants*. MA thesis, UCLA.
Archangeli, D. 1984. *Underspecification in Yawelmani Phonology and Morphology*. PhD thesis, MIT. Distributed by MITWPL.
Archangeli, D. 1985. Yokuts harmony: evidence for coplanar representation in nonlinear phonology. *Linguistic Inquiry* 16, 335-372.
Archangeli, D 1988. Aspects of Underspecification Theory. *Phonology* 5, 183-207.
Archangeli, D. & D. Pulleyblank. 1994. *Grounded Phonology*. Cambridge, MA : MIT Press.
Arnott, D. W. 1970. *The nominal and verbal systems of Fula*. Oxford: Clarendon Press.
Aslin, R.N. 1987. Visual and auditory development in infancy. In J.D. Osofsky (ed.), *Handbook of infant development* (2nd ed., pp. 5-97). New York: Wiley.
Berschin, H., J. Felixberger & H. Goebl. 1978. *Französische Sprachgeschichte*. Munich: Max Hueber Verlag.
Borgstrøm, C. Hj. 1937. The Dialect of Barra in the Outer Hebrides. *Norsk Tidsskrift for Sprogvidenskap* 8, 71-242.
Borgstrøm, C. Hj. 1940. A Linguistic Survey of the Gaelic Dialects of Scotland, Vol. 1: The Dialects of the Outer Hebrides. *Norsk Tidsskrift for Sprogvidenskap*, Suppl. Bind 1, Oslo.
Bird, S. 1991. Feature structures and indices. *Phonology 8*, 137-144.
Bruce, G. 1970. Diphthongization in the Malmö dialect. *Working Papers in Linguistics, Lund University* 3, 1-20.

Burton, M.W., S.E. Blumstein & K.N. Stevens. 1992. A phonetic analysis of prenasalized stops in Moru. *Journal of Phonetics* 20, 127-142.

Calabrese, A. 1993. Marking Statements, Complexity & Simplification Procedures. to appear in *Linguistic Inquiry*.

Campbell, L. 1971. *Historical linguistics and Quichean linguistic prehistory*. PhD thesis, UCLA.

Carnie, A. 1994 *Whence Sonority?* (this volume)

Chambers, W.W. & J.R. Wilkie. 1970. *A short history of the German language*. London: Methuen.

Cho, Y-Y. 1991. On the Universality of the Coronal Articulator. In: C Paradis & J-F Prunet (eds) *The Special Status of Coronals: Internal and External Evidence*. 159-179. San Diego: Academic Press.

Chomsky, N.A. & M. Halle (1968) *The Sound Patterns of English*. New York: Harper & Row.

Clements, G.N. 1985. The Geometry of Phonological Features. *Phonology Yearbook* 2, 225-252.

Clements, G N. 1986. Syllabification and Epenthesis in the Barra Dialect of Gaelic. In: K Bogers, H van der Hulst & M Mous (eds) *The Phonological Representation of Suprasegmentals: Studies Offered to John M Stewart on his 60th Birthday*, 317-336. Dordrecht: Foris.

Clements, G.N. 1991 Place of Articulation in Consonants and Vowels: a Unified Theory. Working Papers of the Cornell Phonetics Laboratory 5. To appear in B. Laks & A. Rialland (eds): *L'architecture et la Géométrie des Représentations Phonologiques*. Paris, Editions du CNRS.

Clements, G. N. & E. V. Hume (to appear) The Internal Organization of Segments. In: J. Goldsmith (ed.) *Handbook of Phonological Theory*. Cambridge, Mass.: Blackwell.

Cole, J. & L. Trigo. 1988. Parasitic Harmony. In H. van der Hulst & N. Smith (eds.) *Features, Segmental Structure and Harmony Processes*, part II. Dordrecht: Foris.

Dayley, J. P. 1985. *Tzutujil Grammar*. University of California Publications in Linguistics.

Dumas, D. 1981. Structure de la diphtongaison québécoise. *Canadian Journal of Linguistics* 26, 1-61.

Eimas, P. D. 1963. The relation between identification and discrimination along speech and nonspeech continua. *Language and Speech* 6, 206-217.

Eimas, P.D., E.R. Siqueland, P. Jusczyk & J. Vigorito. 1971. Speech perception in infants. *Science* 171, 303-306.

Elert, C.-C. 1981. *Ljud och ord i svenskan 2*. Umeå: University of Umeå.

Emeneau, M.B. 1946. The Nasal Phonemes of Sanskrit. *Language* 22, 86-93.

Gamble, G. 1978. *Wikchamni Grammar*. Berkeley: University of California Press.

Garnes, S. 1974. *Quantity in Icelandic: Production and Perception*. PhD thesis, Ohio State U. [Published in 1976 as Hamburger Phonetische Beiträge 18, Buske Verlag, Hamburg.]

Goldsmith, J. 1976. *Autosegmental Phonology*. PhD Diss. MIT. Distributed by MITWPL, Cambridge, Mass.

Halle, M. 1993. *Feature Geometry and Feature Spreading. to appear, Linguistic Inquiry*.

Halle, M. & B. Vaux. 1994. *Feature Spreading and Vocalic Place*. ms, MIT & Harvard U.

Hammond, M. 1988. On Deriving the Well-Formedness Condition. *Linguistic Inquiry* 19, 319-325.

Haugen, E. 1958. The Phonemics of Modern Icelandic. *Language* 34, 55-88.

Hayes, B. 1990. Diphthongization and Coindexation. *Phonology* 7, 31-71.

Itô, J. 1984. Melodic Dissimilation in Ainu. *Linguistic Inquiry* 15, 505-513.

Itô, J. 1986. Syllable theory in prosodic phonology. PhD thesis, UMass/Amherst. Amherst: GLSA.
Kaufman, T. 1971. *Tzeltal Phonology and Morphology*. Berkeley: University of California Press.
Kaun, A. 1993. Coronal Underspecification and Theories of Underspecification. In: D.Silverman & R. Kirchner (eds), *UCLA Working Papers in Phonology* 3.
Keating, P. 1985. CV phonology, experimental phonetics, and coarticulation. *UCLA Working Papers in Phonetics* 62, 1-13.
Keating, P. 1990. Phonetic representations in a generative grammar. *Journal of Phonetics* 18, 321-334.
Kenstowicz, M. & J. Ruback. 1987. *The phonology of syllabic nuclei in Slovak*. Language 63, 463-497.
Kim, C.-W. 1970. A Theory of Aspiration. *Phonetica* 21, 107-116.
Kingston, J. 1985. *The phonetics and phonology of the timing of oral and glottal events*. PhD thesis, UC Berkeley.
Kingston, J. 1990. Articulatory Binding. In: J. Kingston & M. Beckman (eds), *Papers in Laboratory Phonology I*. Cambridge, UK: Cambridge UP.
Kiparsky, P. 1968. Metrics and morphophonemics in the *Kalevala*. In: C. Gribble (ed.) *Studies presented to Professor Roman Jakobson by his students*. Cambridge, Mass.: Slavica.
Kirk, P.L. 1966. *Proto-Mazatec Phonology*. PhD thesis, University of Washington.
Kučera, H. 1961. *The phonology of Czech*. The Hague: Mouton.
Kuhl, P.K. 1987. Perception of speech and sound in early infancy. In P. Salapatek & L. Cohen (eds.), *Handbook of infant perception* (Vol. 2, 275-382). New York: Academic Press.
Kuhl, P. & J.D. Miller. 1973. Speech perception by the chinchilla: Voiced-voiceless distinction in alveolar plosive consonants. *Science* 190, 69-72.
Ladefoged, P. 1975/1993 (3rd edn). *A Course in Phonetics*. Orlando, FL: Harcourt Brace Jovanovich.
Lombardi, L. 1990. The Nonlinear Organization of the Affricate. *NLLT* 8, 375-425.
McCarthy, J. J. 1988. Feature Geometry and Dependency: A Review. *Phonetica* 43, 84-108.
McCarthy, J. J. & A. Taub. 1992. Review of Paradis & Prunet (1991). *Phonology* 9, 363-370.
McCawley, J. 1973. On the role of notation in generative phonology. In: M. Gross, M. Halle & M.-P. Schützenberger (eds.), *The formal analysis of natural languages*. The Hague: Mouton.
Maddieson, I. 1984. *Patterns of Sounds*. Cambridge, UK: Cambridge University Press.
Mascaró, J. 1983. *Phonological levels and assimilatory processes*. Paper presented at the 1983 GLOW Colloquium, York.
Miller, W. 1965. *Acoma Grammar and Texts*. University of California Publications in Linguistics 40.
Mohanan, K. P. 1983. *The Structure of the Melody*. Ms, MIT & National University of Singapore.
Ní Chiosáin, M. 1994. *Against non-constituent spreading: Barra Gaelic vowel copy*. Paper presented at WCCFL XIII, San Diego.
Ní Chiosáin, M. & J. Padgett. 1993. *Inherent V-Place*. UC Santa Cruz Working Papers in Linguistics.
Odden, D. 1991. Vowel Geometry. *Phonology* 8, 261-289.

Otero, C. 1988. From Latin to Romance: the vowel systems. In C. Duncan-Rose & T. Vennemann (eds) *On Language: rhetorica, phonologica, syntactica*. London: Routledge.
Paradis, C. & J.-F. Prunet. 1989. Markedness and Coronal Structure. In: J. Carter & R.-M. Déchaine (eds), *Proceedings of NELS 19*. Amherst, MA: GLSA.
Paradis, C. & Prunet, J.-F. 1989. On Coronal Transparency. *Phonology* 6, 317-348.
Paradis, C. & Prunet, J.-F. 1990. On Explaining some OCP Violations. *Linguistic Inquiry* 21, 456-466.
Paradis, C. & J.-F. Prunet 1991. *The Special Status of Coronals: Internal and External Evidence*. San Diego: Academic Press.
Penzl, H. 1969. *Geschichtliche deutsche Lautlehre*. Munich: Max Hueber Verlag.
Petursson, M. 1972. La preaspiration in islandais moderne. Examen de sa realisation phonetique chez deux sujects. *Studia Linguistica* 26, 61-80.
Poser, W. J. 1982. Phonological Representations and Action-at-a-Distance. In: H van der Hulst & N Smith (eds) *The Structure of Phonological Representations, Part II*. Dordrecht: Foris.
Rice, K. & P. Avery. 1989. On the Interaction between Sonorancy and Voicing. In: *Toronto Working Papers in Linguistics*, vol. 10.
Rice, K. & P. Avery. 1991. *Segmental Complexity and the Structure of Inventories*. Paper presented at the GLOW workshop on the acquisition of phonology.
Rischel, J. 1968. Diphthongization in Faroese. *Acta Linguistica Hafniensia* 11, 89-118.
Sagey, E. 1986. *The Representation of Features and Relations in Non-Linear Phonology*. PhD Dissertation, MIT. Distributed by MITWPL, Cambridge, MA.
Sagey, E. 1987. *Non-Constituent Spreading in Barra Gaelic*. Ms, UC Irvine.
Sagey, E. 1988. On the Ill-Formedness of Crossing Association Lines. *Linguistic Inquiry* 19, 109-118.
Sapir, E. 1930. Southern Paiute, a Shoshonean Language. *Proceedings of the American Academy of Arts and Sciences* 65, 1-296. [Reprinted in: The Collected Works of Edward Sapir, vol. 10 (W. Bright, editor). 1992. Berlin: de Gruyter.]
Schein, B. & Steriade, D. 1986. On Geminates. *Linguistic Inquiry* 17, 691-744.
Schmalstieg, W.R. 1964. The phonemes of the Old Prussian enchiridion. *Word* 20, 211-221.
Shaw, P. A. 1991. Consonant Harmony Systems: The Special Status of Coronal Harmony. In: C Paradis & J-F Prunet (eds) *The Special Status of Coronals: Internal and External Evidence*. 125-157. San Diego: Academic Press.
Sinex, D.G. & L.P McDonald. 1989. Synchronized discharge rate representation of voice-onset time in the chinchilla auditory nerve. *J. Acoust. Soc. Am.* 85, 1995-2004.
Slocum, M. C. 1948. *Tzeltal (Mayan) Noun and Verb Morphology*. IJAL 14, 77-86.
Stemberger, J. 1993. Glottal Transparency. *Phonology* 10, 107-138.
Spajic, S. 1993. *An acoustic investigation of Toda rhotics*. MA thesis, UCLA.
Steriade, D. 1987a. *On Class Nodes*. ms, MIT.
Steriade, D. 1987b. Redundant Values. In: A Bosch et al (eds) *Papers from the CLS Parasession on Autosegmental and Metrical Phonology*. 339-362. Chicago, IL: Chicago Linguistic Society.
Steriade, D. 1987c. Locality Conditions and Feature Geometry. *NELS 18*, 595-617. Amherst: GLSA.
Steriade, D. 1988. Greek accent: a case for preserving structure. *Linguistic Inquiry* 19, 271-314.
Steriade, D. 1992a. *Complex onsets as single segments: the Mazateco pattern*. ms, UCLA.
Steriade, D. 1992b. *Segments, contours and clusters*. ms, UCLA.

Steriade, D. 1993a. Underspecification and Markedness. To appear in: J. Goldsmith (ed.) *Handbook of Phonological Theory*. Cambridge, MA: Blackwell.
Steriade, D. 1993b. *Positional neutralization*. Talk presented at NELS 24, Amherst.
Steriade, D. 1994. *Positional neutralization and the representation of contrast*. Talk presented at MIT.
Thráinsson, Hoskuldur. 1978. On the Phonology of Icelandic Preaspiration. *Nordic Journal of Linguistics* 1, 3-54.
Trubetzkoy, N. 1939. *Grundzüge der Phonologie*. Göttingen: Vandenhoeck & Ruprecht. (English Translation, 1969, University of California Press, Berkeley)
Wescott, R. W. 1965. Review of J.M.C. Thomas: Le Parler Ngbaka de Bokanga: Phonologie, morphologie, syntaxe. *Language* 41, 346-347.
Whitney, W. D. 1889. *Sanskrit Grammar*. Cambridge, MA: Harvard UP.
Woodbury, A. 1981. *Study of the Chevak dialect of Central Yup'ik Eskimo*. PhD thesis, UC Berkeley.

Department of Linguistics & Philosophy
20D-219
MIT
Cambridge, MA 01239

cphill@mit.edu